

# **SPECIFICATION**

Docket No. 0635MH-40874

TO ALL WHOM IT MAY CONCERN:

BE IT KNOWN that we, Frederick S. M. Herz, David C. Parks, and Sampath Kannan, residing in the state of Pennsylvania, Paul Labys, residing in the state of Utah, and Jason Eisner, residing in the state of New York, have invented new and useful improvements in a

## **SECURE DATA INTERCHANGE**

of which the following is a specification:

## CROSS REFERENCE TO RELATED APPLICATIONS

1 The present application claims the benefit of United States Provisional  
2 Application No. 60/161,640, filed October 29, 1999, titled SECURE DATA  
3 INTERCHANGE, and Provisional Application No. 60/206,538, filed May 23, 1999,  
4 titled SECURE DATA INTERCHANGE, both of which are hereby incorporated by  
5 reference.

## BACKGROUND OF THE INVENTION

### 1. Field of the Invention:

7 The Secure Data Interchange invention describes a system to allow a privacy-  
8 protected market for data exchange between multiple self-interested parties. The system  
9 presents a general infrastructure for the exchange of information within a safe privacy-  
10 protected environment, between multiple self-interested parties. We propose a central data  
11 warehouse that maintains data submitted by different users, and executes queries and  
12 programs on the data. Rules are associated with data that define how the data can be used  
13 and queried, to allow agents that submit data to maintain absolute control over its use. SDI  
14 acts as a trusted-intermediary to all parties, and implements an internal market for queries on  
15 the information, allowing agents to specify prices for data access. Furthermore, SDI  
16 supports complex queries such as collaborative filtering, that can provide a querying agent  
17 with a one-time benefit of data access but without long-term access to the data that was used  
18 to compute valuable results.

19 The invention relates to systems that provide personalized information, profiling,  
20 automated matchmaking and information exchange, providing a framework that protects  
21 privacy and allows information collection and profiling within a carefully controlled  
22 environment. Although the marginal cost of data duplication is small, there are hidden costs  
23 associated with data, for example because of privacy concerns, and data can be a valuable  
24 resource in many problems. In business-to-consumer (B2C) applications Secure Data  
25 Interchange addresses the direct conflict between the goal of personalization and the need  
26 for privacy, preventing the exchange and collection of information without knowledge and

1 consent. In business-to-business e-commerce applications (B2B) SDI allows vendors to  
2 provide sensitive and valuable information, for example about business needs and customer  
3 bases, in a secure environment that controls access and leverages value.  
4

5 2. Description of the Prior Art:

6 The invention of Secure Data Interchange relates to a wide-range of application  
7 domains, all of which are characterized by a need to exchange information in a privacy-  
8 protected and carefully controlled market-based environment.

9 As a key application we suggest a system for personalized information delivery in a  
10 networked environment, in which the SDI-proxy can fact as a local filter on information,  
11 based on what it knows about a user's preferences and methods for filtering pushed by the  
12 provider of content. The system allows collaborative filtering through information that is  
13 provided to the central data warehouse, but never released directly to other agents;  
14 collaborative filtering methods are computed in the central SDI data warehouse. Further  
15 motivation is provided with reference to some electronic commerce applications, that we  
16 describe in (A) business-to-consumer and (B) business-to-business e-commerce  
17 applications.

18 In addition to applications within commerce, the system of Secure Data Interchanges  
19 is central to developing many other new products. Examples include the formation of "self-  
20 help" groups between a set of individuals with common interests, and applications to  
21 personal information delivery systems, e.g. for educational and informational purposes.  
22

23 A. Business-to-Consumer (B2C) Electronic Commerce.  
24

25 The recent explosion of electronic commerce, in particular Internet-based individual-to-  
26 business electronic commerce, presents new opportunities for automated personalized  
27 information delivery and the automated customization of products and services. This type of  
28 personalization is very valuable to vendors because it can increase sales volumes, enable  
29 cross-selling and up-selling of goods and services, and allows vendors to price products

1 dynamically based on information about the preferences and goals of customers.  
2 Personalization is also useful to customers when it correctly identifies the requirements and  
3 preferences of a customer, because it can reduce search cost and enhance the “shopping  
4 experience”. Perhaps a customer can find the good or service (i.e. desirable  
5 price/quality/feature tradeoff) that he/she wants more quickly than without personalization,  
6 or receive information about an interesting new product or service that he/she did not know  
7 about.

8         The basis for these new services is that Internet-based “shop fronts” can be  
9 individualized on a per-customer basis, dynamically and in real-time. Traditional main-street  
10 shops must offer the same store layout to every customer, because the layout is physical,  
11 although some level of personalized service can be achieved through well-trained sales  
12 assistants, that act as a “guide” for a customer within a store. On-line “shop fronts” are  
13 virtual, and configurable at negligible cost to the customer or the vendor, assuming that  
14 computation is cheap and fast.

15         Furthermore, Internet-based electronic commerce can allow business to collect vast  
16 amounts of consumer information, because customers interact through a computer-based  
17 interface. Customers can be monitored as they browse a Web site for products and services.  
18 Information such as the search-terms that users enter into a search engine, the links that  
19 users follow, and the length of time spent on each page, can all provide an insight into the  
20 current goal of a customer, i.e. the type of product that he/she wants. When combined across  
21 different sessions, and with similar information about the browsing and purchasing habits of  
22 other customers, the information can be folded into a long-term view of the preferences and  
23 needs of a customer.

24         Moreover, new network connectivity enables different vendors to exchange profiles  
25 for common customers, either statically or dynamically, in order to build broad and detailed  
26 profiles across vendor domains. There exist many potentially powerful synergies between  
27 the data sets that are collected by different vendors, that can be leveraged to provide  
28 appropriate services and products to customers. When analyzed with the proper statistical  
29 tools these data sets can reveal fundamental patterns in the behavior of users, and enable a  
30 vendor to provide appropriate information to a user. Furthermore, access to user-profiles



collected by other vendors can enable vendors to provide focused information delivery to first-time users, and also cross-market services with other appropriate vendors.

Providing user profile information within a carefully controlled environment can benefit vendors and users:

- Vendors would find benefit in sharing data with other vendors; this would deepen their understanding of their customers' behaviors and preferences, especially if some customers were traceable across several data sets.
- Users would benefit from sharing data with other users. This is already evident in the popularity of news groups and web discussion pages catering to individuals with shared interests. By learning what other people with similar tastes and preferences have discovered and enjoyed, a user can sidestep information overload in the search for personally satisfying information.
- Vendors can benefit from receiving data about users. An obvious example would be in the use of collaborative filtering for the marketing of targeted promotions; rather than being deluged with coupons and advertisements that are of absolutely no interest, a user would benefit by being presented with advertising that is highly relevant. In the process, the vendor would increase advertising response rates, boosting overall efficiency.
- Users can receive benefits from providing information to vendors. Personalization of content at vendors' web pages, and well-focused banner advertisements at other web sites that they visit.

The problem is that a user wants controlled personalization, in the sense that it might not be desirable for information about every on-line transaction that a user performs, every on-line document that a user reads, and every web page that a user visits, and demographic information, to be available to every business that the user interacts with, in the virtual and physical world.

#### **A.1 Focused Banner Advertising/Content provision**

Internet-based media sites have followed preceding formats in generating revenue from advertising, with content to users often provided free-of-charge. The business model is

1 similar to that in newspapers, magazines, and television, where circulation and  
2 audience/readership demographics are used to drive revenue. Electronic media presents new  
3 opportunities for media-based business: for example multimedia techniques and  
4 interactivity, personalized delivery of information, and personalized targeting of  
5 advertising.

6 The problem – as before, is to acquire and leverage information about the  
7 preferences and interests of a user, within a system that protects user privacy (i.e. controls  
8 the collection and exchange of information about users, and controls the use that is made of  
9 that information). A further problem is to extrapolate information from a large corpus of  
10 data about an individual user.

## 11 12 **A.2 Mailing Lists**

13  
14 As another example, suppose that business A requests a list of individuals that meet a  
15 particular criteria. Consumer B meets the criteria, but is only listed for business A if A also  
16 meets criteria specified by B, for example if A will provide information about new products  
17 and services that are interesting to B. In an application to the profiling of users on-line, the  
18 problem is that users want to receive the benefits of targeted products and advertisements,  
19 but want to avoid the abuse of profile information and control vendors' access to that  
20 information.

## 21 22 **B. Business-to-Business (B2B) Electronic Commerce**

23  
24 The Internet provides businesses with network connectivity with other business, both  
25 competitors and partners. This connectivity allows businesses to exchange information  
26 about customers (dynamically or statically), in order to identify potential new customers,  
27 build better profiles for existing customers, and up-sell/cross-sell products and services in  
28 real-time. The problem with this exchange of information (that can include swaps, sells, and  
29 rental access) is that businesses need to (a) protect the privacy of their customers; (b)  
30 prevent information release to competitors, either directly or through third-parties.

### B.1 Privacy-Protected Identification of Synergies/Matches

There are many scenarios where autonomous agents would like to be informed of matches under conditions of mutual consent, but without information leakage to any agent if any one of the agents declines the match. Consider two vendors, A and B, and suppose the vendors seek strategic partnerships with other vendors that have appropriate skills and goals. However, vendor A does not want to broadcast to all vendors its need for a business partner or a new alliance, instead vendor A wants to be introduced to another vendor with the right mix of capabilities; similarly for vendor B. What is required is a system that only introduces vendor A to vendor B, and perhaps anonymously at first, if both vendors consent to the introduction. The problem is to provide information that enables matches, without allowing bad matches and abuse of information – i.e. within an environment of secure data interchange.

## B.2 Credential-based Introductions, Contracting and Messaging-systems.

There are many situations where individual parties, for example individuals or businesses, require introductions to credentialed individuals and/or businesses, with the aim of building a new relationship or making a new contract. Consider for example business associations, where credentials about non-bankruptcy, and no previous attempts to defraud could be important. Consider social introductions, where individuals might be concerned about past criminal activities of new contacts. In the domain of automobiles, we could consider a system that identifies other automobiles in the physical location of a vehicle that have recently been involved in an accident. The problem is to manage certificates within a system where users can maintain multiple identities, and to protect the release of certificates without suitable provisions for terms-of-use and criteria for request.

## SUMMARY OF THE INVENTION

1           The above problems are solved, and a technical advance achieved, by the system of  
2 Secure Data Interchange. The Secure Data Interchange system enables information about  
3 bilateral and multilateral interactions between multiple persistent parties to be exchanged  
4 and leveraged within an environment that uses a combination of techniques to control access  
5 to information, release of information, and matching of information back to parties.

6           The system of Secure Data Interchange (SDI) provides a trusted server containing  
7 a large database of information that is owned by its providers. Each data record has an  
8 associated price rule, that controls access to data. The pricing model allows a data owner  
9 to specify a price for different types and amounts of information access, and whether the  
10 identity of the information owner is required, and the system of SDI computes a  
11 composite price for a query based on aggregated prices for a query over a number of  
12 different data owners, with an internal market that favors low priced data. The pricing  
13 model allows discounts based on certificates of a requesting agent, and as a special case  
14 implements the standard capability-based access control systems, where information is  
15 provided to users with appropriate permissions (i.e. with zero and infinite prices). In  
16 addition, the system of Secure Data Interchange allows data to be submitted with a level  
17 of random perturbation (noise), to provide added privacy protection, or alternatively  
18 allow an agent to specify in conditions under which additional noise should be added to  
19 data. A query is priced before execution, to allow an agent to decide whether or not to  
20 execute a query, and select between alternative types of queries. Binding price quotes are  
21 provided to querying agents, and queries can be scaled to meet a budget.

22           Data owners can submit data to the central data warehouse with different degrees  
23 of identification, for example anonymously, pseudonymously, or with a true identity. For  
24 example, in the case of data that represents a user's profile information, for example  
25 information about the interests of a user, a user might prefer to use a number of different  
26 pseudonyms for different types of activities that he/she likes to engage in online. A user  
27 might maintain a number of different aliases within the database, for example to represent  
28 different types of things he/she likes to do which have little bearing on each other.



1 manages a user's interactions with the central SDI data warehouse, i.e. providing profile  
2 information and controlling profile access. The client-side SDI proxy for an agent that  
3 represents an individual browsing the Internet can manage that user's profiles in  
4 interactions with other agents, for example representing vendors and content providers.  
5 The client-side SDI proxy can also handle decisions about what types of information to  
6 submit to the server, and manages query execution on behalf of the agent. The client-side  
7 SDI proxy agent can also push information about a user's on-line activities to the central  
8 SDI data-warehouse in real time. This enables a system of "time-of-purchase-  
9 competition" system, in which a user can request competitive counteroffers from other  
10 vendors before making a purchase.

11 The system addresses the fundamental conflict that exists between rights of privacy  
12 and efficiency gains from better bilateral exchange of profile/preference information. SDI as  
13 applied to B2C e-commerce allows consumers to receive targeted information about  
14 products and services, but without the loss-of-privacy that can easily occur in the current on-  
15 line profiling "free-for-all". The cookie technology provided by Netscape to supported  
16 personalized sessions with a single vendor on-line has been used by advertising network  
17 providers such as DoubleClick to track users across multiple sites, often without either the  
18 consent or knowledge of that individual [New York Times, Feb 7, 2000].

19 In describing the system of secure data interchange we claim the following novel  
20 technical ideas:

- 21 (a) Agents can associate price-rules with information that is placed in the central  
22 database, and retain absolute control and ownership over all uses of that data.  
23 As a special case of price rules, the system supports access based on certified  
24 properties of querying agents (with zero and infinite prices). The pricing  
25 model allows information providing agents to receive direct value for data,  
26 and allows agents that request access to information to receive a price before  
27 a query is executed, and make appropriate decisions about what type of  
28 queries to execute.
- 29 (b) A number of novel techniques are proposed to allow data processing within  
30 the data warehouse without releasing too much information to an agent.  
31 General programs, for example collaborative filtering techniques, can be



1 about a user's profile. In a simple form, the vendor provides complete information about  
2 its services, and a method to display them to the user based on his/her local profile. When  
3 describing the application of SDI to electronic commerce we also describe methods to  
4 implement necessary ancillary systems that are essential to supporting full e-commerce  
5 functionality within an identity-protected system, such as systems for pseudonymous  
6 payments and physical mailing of products.

7 Collaborative filtering based on profiling information from multiple users is  
8 supported within the central SDI data warehouse, but within a system of economic  
9 incentives, where users provide profile information in return for receiving payments from  
10 vendors for that information. This allows broad network-wide information to be used for  
11 profiling, in addition to deep vendor-specific information.

09699098 "102700



## BRIEF DESCRIPTION OF THE DRAWINGS

1        The novel features believed characteristic of the invention are set forth in the  
2        appended claims. The invention itself however, as well as a preferred mode of use, further  
3        objects and advantages thereof, will best be understood by reference to the following  
4        detailed description of an illustrative embodiment when read in conjunction with the  
5        accompanying drawings, wherein:

6        Figures 1 through 21 illustrate various parts and embodiments of the invention.

002207-36066960

## DESCRIPTION OF THE PREFERRED EMBODIMENT

### 1. Introduction

The invention of Secure Data Interchange (SDI) describes a general infrastructure for the exchange of information within a controlled environment. We propose a central data warehouse that maintains data submitted by different users, and executes queries and programs on the data. Rules are associated with data that define how the data can be used and queried. As such, the system of SDI prevents the exchange and collection of information without knowledge and consent. The system allows for payments to be received by the providers of information, in return for data access. In application, the invention enables new systems for the delivery of personalized information, profiling and automated matchmaking and information exchange, all within a framework that protects privacy and maintains data security. SDI supports the collection and exchange of information between, and relating to, autonomous (and possibly self-interested) agents within a distributed environment.

#### 1.1 Definitions.

**Agent.** An agent in SDI is a any party that wants to sell or give away data to other parties, or buy or receive data, or in general both provide and receive data. Agents may represent any party with individual goals, autonomy of control, and a persistent identity. Examples in business-to-consumer e-commerce include business such as newspapers, book stores and travel companies, that wish to receive data about the profiles and buying habits of users so that they can personalize the information, products and services that are sold to users. We assume that agents are autonomous from the system of Secure Data Interchange, and follow actions consistent with their preferences, abilities and resources. Similarly, an on-line consumer is an agent that wishes to provide data about its preferences and buying habits to vendors and other consumers, in return for well targeted products and financial reward.

**Agent Computers.** Agents are represented in the system of secure data interchange with dedicated computational resources, agent computers, with permanent memory, processing power, and network connectivity. For example, an agent computer might be a consumer's home PC that will act as a client machine in interactions with the central SDI data server and



1       The identity that an agent assumes with a profile matters not only to the ability of an  
2 agent to control the agents that can send it messages, but also its ability to control the  
3 amount of information that can be exchanged about the agent in the open marketplace, out  
4 of the agent's control. A basic premise in the system of Secure Data Interchange is to keep  
5 control of data. We provide methods that allow an agent to release data, or performing  
6 processing on data, to one agent but prevent that agent from selling the data on to another  
7 agent with which agent A also interacts. For example, one key technique is for agent A to  
8 use a unique pseudonym with every agent that it interacts. This allows an agent (so long as it  
9 is careful not to release other identifying information) to release profile information to other  
10 agents without losing the value of that information, because there can be no secondary  
11 market in the agent's profile. If agent A provides information to agent B, then agent B  
12 cannot pass that information onto another agent C and have it still linked to agent A because  
13 agent C does not know the identity of agent A among the agents that it interacts with.

14       An agent's profile contains any and all data that an agent might wish to exchange  
15 with another agent when it assumes a particular identity. However, just because the data is  
16 in the profile it does not mean that it is available to another agent. Possible information in  
17 the profile of an on-line consumer includes: transactions that it has performed with other  
18 agents; information that relates to its true identity (e.g. salary range or education level);  
19 information provided by the agent (such as its preferences for a particular type of product,  
20 etc.); and other information that has been compiled based on observing the behavior of the  
21 agent (e.g. physical location for a mobile user, such as a user in a vehicle, or trace of recent  
22 web pages visited for a user that is browsing the Internet.)

23 **SDI-Proxy:** An SDI-proxy refers to the software that runs on top of an agent's computer  
24 device, and configures that device for Secure Data Interchange. The proxy intermediates  
25 interactions between pairs of agents, and also intermediates transactions between agents and  
26 the SDI data server. For example, an on-line consumer might define a profile-management  
27 policy at the client-side SDI proxy that automatically configures the agent's profile and  
28 identity when as the agent interacts with other agents. The profile-management policy  
29 implements an appropriate policy to select the user's profile and identity on the basis of the  
30 information that is available about the other agent (for example from certificates).

## 1.2 System Architecture

The basic architecture for SDI is a system of agent computers, connected via a network (Internet, wireless, or otherwise) to other agent computers, and with a central SDI shared data warehouse. In Figure 1 we illustrate the top-level architecture of Secure Data Interchange, the networked system of agent computers and a centralized server computer that acts as a repository for data, rules and code. This is called the SDI data warehouse. Each user is associated with an agent computer, and in general users can be individuals, groups of individuals, or companies. In its most general form, the system of Secure Data Interchange is for a system of multiple autonomous agent computers, involved in multilateral communication. We restrict the system description to bilateral communication between agents, without loss of generality because any multilateral (multicast or broadcast) can be implemented as a set of bilateral communications. In any bilateral communication there are two parties, the sender and the receiver (and the parties can dynamically change over a communication session).

SDI proxy agents and profile management policies on agent computers mediate bilateral agent interactions. Agents define profile-management policies that are implemented on agent computers and determine appropriate profiles and identities for an agent in interactions. Every time an agent initiates a new interaction, new information is available about the agent, that can be compiled by the agent's own computer device, and also by the computer device of the agent with which it interacts. Careful profile and identity management provides an agent with absolute control over the ability of other agents to profile the agent and exchange information about the agent, for example the agent can use a unique pseudonymous identifier with every other agent. However, the core of the SDI invention is that we encourage agents to exchange information, by providing a secure central data interchange for that purpose. Agents can submit profile information, and other data, to the central data warehouse and make it available for particular types of data mining by other agents, and receive financial reward for providing data. Products and services can also be targeted for agent A on that agent's computer, without the provider agent receiving information about agent A's profile. For example, vendor B might provide generic information about its products to agent A, and a personalization rule that it has constructed

1 from data mining the central SDI data warehouse. The agent computer of agent A can then  
2 use its local and private profile information to decide what products to target to the agent,  
3 based on the instructions that it receives from agent B.

4 The arrows indicate possible flows of information between the different elements  
5 within the system. We allow agents to communicate directly, via bilateral or multilateral  
6 communication, and also indirectly via the central data warehouse. Information can be  
7 submitted to the central data warehouse, for example new data and new rules, or queries  
8 that the data warehouse will execute. The data warehouse returns information to agents, for  
9 example the results of a query. The Secure Data Interchange invention is independent of the  
10 implementation details of the communication platform. We assume that the system of SDI is  
11 built on top of a secure communications platform, for example via a SSL-encrypted TCP/IP  
12 session in an application to the Internet. Furthermore, we draw on cryptographic techniques  
13 known in the art for identity management, and additional techniques to support  
14 pseudonymous identities within a particular network protocol, e.g. the stripping of sender  
15 URL addresses from HTTP packets on the Internet.

16 In overview, the top-level claim in the system of Secure Data Interchange, of a  
17 privacy-protected market-based system for data exchange between self-interested parties, is  
18 constructed from the following core components:

- 19 ○ A Central Data Warehouse, with associated price-rules and constraints submitted by  
20 the owners of data. Architectural variations allow virtual links to data in the central  
21 data warehouse, with data physically located in distributed locations.
- 22 ○ A Query-execution and Price module, which executes queries on the data within  
23 constraints on the types of queries that can be performed, and computes the price of  
24 queries, collecting payment.

25 and the following optional components:

- 26 ○ Distributed data management, via client-side SDI “proxy” agents, that intermediate  
27 the interactions between agents, and manage data provided to the central SDI data  
28 warehouse. Distributed data management includes methods to manage an agent’s  
29 identity in interactions with other agents, for example via pseudonymous and  
30 anonymous interactions.

- 1       ○ Distributed query-execution, via client-side SDI query-execution modules, that
- 2       allow data that is physically stored in distributed locations to be queried on
- 3       distributed nodes, without an explicit release of the data.
- 4       ○ Data perturbation methods to augment data values with randomized noise, to allow
- 5       queries to tradeoff price and quality, and to allow a user to protect his/her identity by
- 6       hiding revealing details of submitted data elements.

7       In describing key applications of the system of Secure Data Interchange, we describe  
8 in later sections additional features that extend the functionality of the invention:

- 9       ○ Community dollars: an extended payment scheme that allows extended forms of
- 10       payment within the system, for example allowing payment to be made in terms of
- 11       discounts in purchases from particular vendors, or in return for agreements to
- 12       purchase a certain number of products over an extended period of time.
- 13       ○ Within an Internet-browser based system, methods to control a user's browsing
- 14       experience via a personalized portal, where personalization is performed via data
- 15       mining techniques executed by the system of SDI on data in the central data
- 16       warehouse.
- 17       ○ Methods to support anonymous and pseudonymous electronic commerce, e.g.
- 18       delivery and payment services.

19       In the next section we provide an overview of each of the central SDI data warehouse,  
20 which forms the core component of Secure Data Interchange. We then add technical  
21 details for specific components, to make our ideas more concrete.

## 22       1.3 Implementation Details

23       In this we describe some of the core technologies known in the art that would be used to  
24 build up an implementation of the SDI system. The technologies span areas in  
25 cryptography, for pseudonym management, digital certificates, payment mechanisms,  
26 etc.; and information theoretic-methods, for example to protect the identity of the  
27 originator of a message by routing messages through local "crowds" of agents.

### 28       1.3.1 Certificate Management

Cryptographic techniques well-know in the art [Chaum 81; Chaum 85; Chaum 91] provide the ability for a certificate to be linked to an agent's identity, and not transferred to other agents. This is the basic functionality required of a certificate management system. For example, a certifying agency can sign the public key of an agent with the private key of the certifying agency, to indicate that the agent satisfies requirements for certification. Another agent can verify the certificate with the public key of the agent and the public key of the certifying agency. The certificate cannot be transferred to another agent unless that agent assumes the same public key. We assume a public key infrastructure to manage this process.

### 1.3.2 Support for Anonymous and Pseudonymous Identities

We have already noted that it is important to provide whatever additional support is required because of the underlying communication infrastructure to protect agents' profile management policies. For example, in the TCP/IP mechanism a message must be stripped of the network address of the originating Internet server, because this can provide information to allow pseudonyms to be linked. Similarly, messages can be routed through a common gateway or random "forwarders" as in the "CROWDS" system [GGMM98; RR98] to provide pseudonymity. Furthermore, other e-commerce functions, such as payment and the anonymous mailing of goods must be supported (see [INSERT A FORWARD REF TO A LATE SECTION OF BOTTOM-LEVEL SDI DOC]).

The ability to embed data within web pages allows client-side processing of information. By embedding profile and location information directly within a web document we can alleviate the bandwidth and computational bottlenecks that can occur at a centralized profile server if profiles are fetched on-the-fly when web pages are downloaded by clients. The origin server (supported by the vendor) requests periodic profile updates from the central SDI server. This duplication of information enables the profile and the page contents to be provided directly from a vendor's server.

There are some potential drawbacks of this approach: (1) the profile information associated with a web page and target objects can be out-of-date; (2) the profile information is available to all clients and proxy servers, not just those that are SDI-enabled; (3) the profile information can be altered. We suggest technical solutions to each of these problems below.





1 In addition, we prevent unauthorized access of embedded profile information  
2 through the encryption of the metadata that is represented within the XML structure of a  
3 web page.

4 Profile information can be encrypted using a hierarchy of keys, so that different  
5 levels of access to the information may be provided according to the access levels of users  
6 and vendors. All users that request web pages from SDI-enabled vendors, whether or not the  
7 user is a member of SDI receive the same profile information. We provide encrypted  
8 profiles to vendors in the 'Profile Update' messages from SDI to vendor servers, so that: (a)  
9 unauthorized agents cannot tamper with the profiles; (b) the profiles cannot be read by  
10 unauthorized agents.

11 The SDI system supplies a private key to trusted SDI client software, that enables  
12 only SDI-enabled clients to access profile information, and only access that information to  
13 the extent permitted by privacy policies of users and vendors. Different levels of encryption  
14 enforce multiple levels of access. Periodically the key pairs are changed to prevent  
15 extended attempts at cryptographic attacks. The SDI system uploads the key that provides  
16 the correct level of access for a user to a user's client, once terms of access and profile  
17 management have been agreed. A client can only access embedded information once  
18 enabled with a relevant key. Finally, profile information is signed with a digital certificate,  
19 to prevent third parties from tampering with profiles for commercial gain.

## 20 2. Core Modules

21

22 Figure 2 illustrates the core modules within the system of Secure Data Interchange. In  
23 this section we briefly describe each module in relation to the other modules, and provide  
24 more details in later sections.

25 The central SDI data warehouse is the core system in SDI, managing data records on  
26 behalf of agents. In combination with the query-execution module, these two modules  
27 implement the privacy-protected market for query-execution, where queries are executed:

- 28 • If the querying agent has the right certificates, as defined by the owner of data
- 29 • (and) If the querying agent pays the cost of performing a query





1 maintain control over access to the data. Price rules can also specify different prices  
2 based on the degree of random perturbation that is performed on data before a query is  
3 executed. This can protect the identity of the owner of information but still allow  
4 valuable information to be provided in response to a query.

5 The system of SDI also expands the possible space of agreements through anonymity  
6 techniques and through random perturbation of data. Agents can provide information  
7 anonymously or pseudonymously, or with their true identities revealed. Furthermore,  
8 agents can charge less for information provided anonymously than for information  
9 provided under a revealed identity. Meta-information associated with data can also  
10 specify whether the information has been certified by a third-party, or whether it has been  
11 randomized slightly before placing in the SDI data warehouse.

12 The data repository is linked to the query-execution module, that performs queries  
13 that are queued to be executed in the pending queries queue. We allow queries to be  
14 general programs, an extension from the simple queries that are found in SQL-based  
15 database languages. The key novelty is that the query-execution module can maintain  
16 intermediate results, for example pointers to records selected as part of intermediate  
17 queries. In a standard database there is more of a separation between data access and data  
18 processing, with data accessed, pulled outside of the database, and then processed. We  
19 allow data to be accessed and processed in the database, with results pushed to clients.  
20 This is has useful privacy properties, because less information is finally provided in  
21 response to a query. It is very useful in the space of secure data interchange because it  
22 expands the set of agreements that can be reached between providers of information and  
23 requestors of information.

24 Data manipulation is via record pointers, which are one-time and anonymous pointers  
25 to data records. We associate a single query (possibly a complex query, i.e. a query  
26 program) with a query session. A record pointer is a temporary identifier for a record that  
27 is selected. The pointer allows persistent queries to be performed on the same record  
28 during a single query session, but is not valid in other sessions. The same data record  
29 might be accessed multiple times during a single session, and with multiple record  
30 pointers when the accesses are independent (i.e. when the querying agent does not know  
31 that it is requesting information of the same data record). Provisional payments are tallied

1 against each record pointer, and relate to the degree of information provided about that  
2 record. For example, consider the problem of finding a set of profiles of a particular type.  
3 As the search is performed over data records in the database the query-execution module  
4 assesses a provisional payment, as records are accessed. However, the requestor of the  
5 information only pays the sum payment charged by the records that are actually selected  
6 as suitable. The owners of the data records that were not selected in the final result  
7 returned to the user receive no payment, because no information about their data was  
8 released and the information was only used in intermediate processing steps. This is  
9 important, because it makes it unnecessary to formulate queries carefully in order to  
10 avoid extra cost because of redundant intermediate steps.

11 Complex queries, including query-programs, with intermediate results and  
12 anonymous record pointers expand the space of data manipulation in a system with  
13 providers of information that are more willing to release information if that information is  
14 never made directly available to another agent except in aggregated system-wide terms.  
15 An agent might be willing to allow submitted data to be used for data mining applications  
16 so long as the data remains in the trusted data repository. The central SDI server can  
17 provide standard types of query-programs, stored in the SDI-methods module, for  
18 example methods to perform collaborative filtering.

19 The query-execution module also performs persistent queries that are submitted by  
20 agents to be executed when particular conditions in the data repository are true. Persistent  
21 queries check for certain information to reside in the data warehouse, and notify a  
22 requesting agent whenever these conditions exist, perhaps automatically making an  
23 action.

24 The pricing module is responsible for pricing queries before execution. The basic  
25 problem is to aggregate the charge over all data records that provide information to form  
26 part of the result of the query, and to estimate the price of a complex query without  
27 performing the query. We suggest a simple top-level query language to allow a user to  
28 price quality-cost tradeoffs in the query that is finally executed. For example, statistical  
29 techniques can be used to compute aggregate statistics without accessing all data records  
30 in the database. To give another example, a query can be performed on the basis of  
31 accurate information or on the basis of slightly randomized information. The pricing



applications, where job offers can contain salaries that are only revealed to applicants with suitable qualifications.

SDI is designed to allow users to submit many different types of information, from many different sources, and for many different purposes. The application is particularly useful when it becomes a *de facto* data repository for lots of different types of information, information that can in fact be *analyzed* for the purpose of extracting patterns and other useful data (e.g. with collaborative filtering techniques). However, we provide three key variations in the following description, which are not necessarily exclusive.

1. A single unified database, with information residing on a central server (or a number of central servers). We describe two possible implementations for such a heterogeneous database, one XML-based and another object-oriented model.
2. Multiple databases, partitioned into “data types”. For example, one SDI database might contain profiles on online customers, while another database might contain information about job offers and job applications. This type of database can be implemented using a standard “indexed-field” representation. Again, the information resides on a central server (or a number of central servers).
3. Distributed data, with some data residing on a central server, and some data physically located on distributed servers (e.g. on the servers of vendors that subscribe to SDI), but with virtual “hyper links” from the central server to give one unified view of the data.

The key over-riding problem is one of *data representation*, the way in which information is encoded. We describe a general infrastructure for information exchange, and do not mean to limit the description to any one type of data. Furthermore, the system of SDI should be *expandable*, so that new data types can be introduced. In this specification the type of information in the Secure Data Interchange shared data warehouse is not constrained in any way, but may include for example user profiles (e.g. preferences, recent purchases, etc.) or business services (e.g. costs for services, service capabilities, etc.).

One approach is to use a single native SDI ontology, which is expanded as necessary. It would be the responsibility of providers of information in alternative forms



to provide “translation services” to convert local data formats into the SDI native data format. A typical technology to support a shared ontological representation is XML (Extended Meta Language), which allows a grammar to be defined for a document, with meaning embedded in tags. The trend towards XML-based applications should facilitate a shared ontology structure, and allow metainformation to be associated with information and describe data. XML allows intelligent integration of data from multiple databases.

Alternatively, we can allow data to be stored in heterogeneous formats across a single unified database, within an object-oriented infrastructure. Each data object has a “wrapper” that controls access, and provides an interface for queries. When a query is executed, the method is invoked, and the result computed with the method and the data. This is perhaps more efficient than the aforementioned approach, in that efficient data formats are query structures are retained.

### 3.1 Data Structure

In this section we describe the structure of the records in the SDI database. The next section describes the interface that allows user agents to submit data and update information.

The data repository has the following key features: *price rules* are associated with information, to control access to information on the basis of information requested and properties about the requesting agent; meta-information to specify additional information about data records, for example has the information been certified by a third party.

We describe the abstract structure of data records in the data repository. The semantics of the data field and the price rules depend on the type of information. All data records, whatever the information that they represent, contain the following elements:

1. Owner ID.

The owner ID is a three-tuple (Public Key, SDI identity code, Remote address).

The public key is provided by an agent that submits data, and is one half of a public/private key pair in a public key based cryptographic infrastructure. The public key is used to provide authentication of the agent, in case it wants to amend the record in the future. The agent can submit a message signed with its public key to prove its identity. An agent can use a different public key for each alias that it maintains within SDI. The SDI identity code is a unique code, generated by SDI for

each agent alias, and provided to an agent to allow the agent to access the eBank and other ancillary SDI services, for example to collect payments received for access to the data. The remote address is an (optional) contact address for the agent that submits the information, for example an email address to a pseudonymizing module that will forward email to the agent under a number of different aliases.

## 2. Record ID.

The record ID is generated by SDI, and is used in the case that a user creates a number of different data records with the same owner ID.

## 3. Data object.

As noted above, we allow data in SDI to be of different types, and a general data object might be represented in an internal coding that is not known to SDI. Such an object must provide an *accessor* function to allow queries to be performed. The query-execution module takes an object and invokes the accessor methods to perform a query. In the most general form, a data object is represented as a three-tuple: (data type, data field, data accessor). The data type specifies what the type of data is, for example is it a user profile or a business rule. The data field specifies the data, and can be a private record that is only accessed by the accessor, which provides an interface to allow queries of the correct format for the data type to be performed.

In a simple special case, with data types designated by SDI centrally and data indexed with fixed fields, a data object is more appropriately represented as an association list between field names and values, e.g. (Field1, Value1), (Field2, Value2), etc..

## 4. Price rules.

The price rules provide the owner of information with control over the type of queries that can be performed on a particular data record. The price rules compute an ask price for a proposed query, on the basis of three pieces of information: what type of information, and what accuracy of information is requested; what certificates can the querying agent present; can the information be provided anonymously or must the identity of the owner of the information be revealed? As special cases, a price \$0 corresponds to "access is possible for free", and a price \$infinity corresponds to "no access is possible". The unit of currency need not be US dollars, but can be any unit



1 In general, a query of a data record can be part of a larger query session, as  
2 discussed in the section on query execution. Within a larger query session, the same data  
3 record might be queried a number of times, for example with a large compound query  
4 split into a number of steps with continual execution contingent on continued correct  
5 responses. E.g. Select records of type A, then from those records select records of type B,  
6 then from those records select records of type C... In this case, because price rules might  
7 be *non-linear* in the amount of information provided, it is important to provide a price-  
8 rule method with information to allow it to track a sequential query.

9 As discussed in more detail in the next section, the system of SDI handles this by  
10 generating temporary and anonymous pointers to data records, that are valid only for a  
11 single session, and allow a price-rule to track sequential queries. The data record pointer  
12 allows the history of queries to be recovered, and a new price to be computed on the basis  
13 of *total information* provided. Similarly, this is important when a query might collect a  
14 lot of information about a data record and then request the identity of the agent that  
15 provided the information to SDI. The cost of revealing an agent's identity (and therefore  
16 allowing a secondary market in its information) might well increase with the amount of  
17 information that has already been released.

18 Similarly, in some cases the same data record might be accessed on multiple  
19 occasions, but independently, such that the environment performing queries does not  
20 know that the data record is the same record from before. E.g., select records of type A  
21 and extract information with rule B, then select records of type C and extract information  
22 with rule D. A single record can be of type A *and* of type C, and therefore be selected for  
23 information extraction in both cases. In this case, each independent sequential query (i.e.  
24 A, B and C, D) has a set of data record pointers, so that the cost of extracting information  
25 is computed independently for any record that has type A and type C.

26 Here is a simple example of a non-linear pricing function, that accounts for  
27 cumulative information that has been provided to a requesting agent. Suppose that any  
28 sequential query receives a temporary and anonymous pointer to a data record, that is  
29 only useful in the current query session. The *data record pointer* allows the price rule to  
30 implement a non-linear pricing rule. For example, consider the query 'what is the value

1 of field A, B and C', split into queries 'what is the value of field A', then 'what is the  
2 value of field B', then 'what is the value of field C'? The pricing rule might state:

3 \$0.1 for any one of A, B or C

4 \$0.2 for any pair of A, B and C

5 \$10 for all of A, B and C.

6 In this case, by tracking the data record pointer with which a data record is accessed, the  
7 price of the first two requests can incur an incremental charge of \$0.1, while the price of  
8 a third request can incur an incremental charge of \$9.80.

9 At a per data-record level the price of a query depends on the response made to  
10 the query. This is important because an affirmative answer to the question 'are you the  
11 president of the united states' carries more information than a negative answer. This has a  
12 slightly undesirable side-effect, in that when estimating the price of a query in the pricing  
13 module, the *estimated price of a query might leak information about the result, even*  
14 *without performing the query*. However, we believe that the benefits of linking price to  
15 information content in a query outweigh this potential loss in value of information. Note  
16 in particular, that in general the ask prices associated with data records are private  
17 information and independently set, and therefore are not very revealing, especially within  
18 a competitive market place.

### 19 3.3 Data Submission/Update Methods

20 In this section we describe the basic methods to register with SDI, submit data, update  
21 data records, and access payments collected by SDI as information is queried. We use  
22 "agent" to refer to the computer system that interacts with the central SDI data  
23 warehouse, submitting data and requesting payment from the eBank. Agents may  
24 represent individuals, vendors, or other self-interested parties.

25 Data records can be submitted under multiple aliases by a single agent, to provide  
26 an additional level of control and flexibility in managing data submitted to SDI. For  
27 example, if an agent represents an individual that is an online consumer, interacting with  
28 different types of vendors, then perhaps the agent will use two aliases: one for while the  
29 consumer is at work, and one for while the consumer is at home. Alternatively, the agent  
30 can maintain a number of identities for different activities, or interests of the user. In the  
31 preferred implementation agents, for example web-browser based client agents, can

1 manage an agent's selection of identities as it submits data to the central SDI warehouse  
2 local to the user. The ability to submit information under multiple identities protects the  
3 ability of an individual to prevent another agent building a complete picture about its  
4 preferences and profile, while still allowing that individual to leverage as much of the  
5 value associated with its information as possible. The SDI proxy agent, situated on user's  
6 client machines, implements this functionality. It is described later in this patent.

7 A standard cryptographic public key /private key infrastructure provides a useful  
8 technique to implement a system in which agents can maintain different aliases. As  
9 proposed in the work of D. Chaum [Chaum 81; Chaum 85; Chaum91] a public key (PK)  
10 and private key (SK) serves a number of purposes. First, the public key acts as an  
11 identifier for the alias, a name. Second, the agent can compute a new private key/public  
12 key pair when it requires a new alias. The agent keeps the private key secure, and this  
13 provides a method to allow the agent to validate its identity, for example by  
14 cryptographically signing a message with the private key. The signature can be verified  
15 with the public key, and the keys can be selected with enough bits to make falsification a  
16 computational impossibility. This infrastructure is outside of the current patent, but  
17 standard in the art.

18 As described in the JANUS/LPWA system [BGGMM 97; BGGMM 98; GGMM  
19 98], it is also possible to associate a public/private key pair with a pseudonymous e-mail  
20 address, to allow information to be pushed to an agent that owns information under its  
21 alias. The system is implemented via pseudonymous proxies which a user agent to poll  
22 and check for new messages.

23 An alias can also be completely *anonymous*, but in this case the value of the data  
24 provided may be less in the internal market place implemented within SDI (in the query  
25 execution module). The convention for an anonymous alias is that the user agent  
26 continues to provide a public key, and use a private key for validation of its identity.  
27 However, in this case the public/private key is one-off and just for this data record, and  
28 no return address is provided. An agent that submits information anonymously can still  
29 recover payments from the eBank.

30 When a data record is first created SDI returns a data ID, so that the agent that  
31 submits information can specify a particular data records in future, in case it creates a







1        object-oriented framework with queries performed by invoking methods that are  
2        associated with data.

3 e. Monitor conditions for the persistent queries, and provide a response to a query if  
4 conditions are satisfied.

5       A central part of the invention of SDI is the method to compute the price of a query as  
6       it is executed. We described the semantics of price rules, which price access to data.  
7       They are defined by the owner of information. The SDI query execution module  
8       implements an internal market for queries, and ensures that queries are executed at  
9       minimal cost to agents. The SDI query execution module is also responsible for  
10      collecting revenue on behalf of owners of information.

We describe this methodology in this section, describing how the total price of a query is computed, as a sum of the price charged by the owners of data records that provide information which contributes to the final response.

14 The query execution module follows a protocol to execute queries:

- ```

15   1. Request_For_Price(Query, Agent_Certificates)
16   2. Estimate Price with call to the Pricing Module
17   3. Price_Quote(Quote_ID)
18   4. Request_Query(Quote_ID)
19   5. Get_Payment()
20   6. Execute the Query
21   7. Report_Answer()

```

22 In step (1) an agent makes a request for a query to be priced, stating the query, and  
23 providing certificates to allow the query to be priced. In step (2) the query execution  
24 module makes a call to the pricing module, and a price for the query is computed  
25 (described in the next section). In step (3) the agent receives a price quote, and can then  
26 decide whether or not to execute the query, and also a query ID. In step (4) the query  
27 execution module receives a request to perform the query, and then in step (5) requests  
28 payment from the agent. When payment is received the query is executed (6), and  
29 appropriate payment is credited to agents that provide information, scaled to make the  
30 budget balance as necessary (in the case of an over or under price quote). Finally, the  
31 response to the query is provided (7).

#### 4.1 Price Rule Semantics

Each data record has an associated price rule that controls the price of accessing information. A data record can contain a number of different pieces of information, and data can be provided to different degrees of accuracy, so the price rule can be quite complex in general. In particular, consider a data record that represents the profile of an individual. The individual might be happy to have information released about some fields, for example its ZIP code, or its recent book purchases, but less happy to have information released about different fields, for example its salary or social security number. Clearly, different pieces of information within a single data record require different prices.

Furthermore, a price rule for the data record as a whole might need to be superadditive across data elements, such that it becomes very expensive to request too much information about data associated with the same user. One of the driving concerns behind the present invention is that at present it is possible for on-line vendors to collect information about a single individual via “cookies” (identifying codes which are left on a user’s client machine), and form a portfolio of information about various activities and preferences of that individual.

While information about an individual might be acceptable in small amounts, in large amounts the same information can soon become unacceptable. Furthermore, even if users release information anonymously in response to a query, if a lot of information is released the identity of the user can be compromised. Every additional piece of independent information that I state about my profile identifies myself a little more clearly, and acts to distinguish me from the profiles of other individuals. We discuss this further in a later section on random data perturbation, which describes how random noise can be added to data to counteract this effect.

The basic idea is that the system of SDI allows a user to associate a price rule with every data record, that computes the price that a user must pay to execute a query over that data. We allow the price to vary, depending on properties about the requesting agent, the amount of information requested, and the level of identification that is required of the owner of the information in responding to the query.

The abstract form of a price rule for a data record is a function:

1 ( Information-request x Certificates x Level-of-identification ) → Price

2 Given a request for information, a set of certificates that a requesting agent can present,  
3 and the level-of-identification that is required of the agent that provides information (i.e.  
4 anonymous, or revealed identity), then a price rule computes an ask price. This is the  
5 price that an agent must pay to execute the query on the data record.

6 However, as we describe in the next section, an agent only actually pays this price  
7 if the information provided is used to compute information in the final response provided  
8 to a query. For example, consider query "Select all data records close to record X". The  
9 query is most simply executed by computing the distance between each record and record  
10 X. Although a negative response from a record that is not close to X provides some  
11 information about that record, the information is not used to compute the information (i.e.  
12 set of records) that is finally provided in response to the query.

13 As in the representation of heterogeneous price rules in SDI, the invention of SDI  
14 allows heterogeneous methods to compute prices for queries. In the most general case,  
15 we allow a price rule to be computed as a price method, which is invoked for a data  
16 record before a query is to be performed. The query-execution and pricing modules  
17 simply invoke the price method, and the data object provided by an agent returns a price  
18 for the query. In this most general version, the role of SDI is limited to providing the  
19 price method with the query, the certificates of the requesting agent, and the level-of-  
20 identification that is required.

21 In a simpler variation the system of SDI can provide a number of default price  
22 rule languages, which allow a user to specify in simple but quite flexible terms a price  
23 schedule for queries. We describe two such rule languages: an uncertainty-based additive  
24 price rule; and a grouped additive price rule which is a simplified version of the general  
25 uncertainty-based rule.

26 The uncertainty-based additive price rule allows a user to adjust the price for a  
27 query based on information about the requesting agent, the amount of information  
28 requested, and whether the information is required anonymously or with a revealed  
29 identity. It does not allow a non-linear coupling across the price of information about  
30 multiple attributes, but instead includes a simple upper-bound on the amount of  
31 information that can be requested. Such a non-linear coupling could be added with simple









- 1 • Low price, if an agent presents a Charge certificate, and requests information with
- 2 a High degree of uncertainty.
- 3 • High price, if an agent presents a Charge certificate, and requests information
- 4 with a Low degree of uncertainty.
- 5 • Infinite” price, if an agent cannot present a Charge certificate.

6 Different price functions can be defined for anonymous and revealed-identity. The  
7 price rule has the following components:

- 8 a) A set of certificates that allow data access for free.
- 9 b) A set of certificates that allow data access for some charge.
- 10 c) A list of price-functions, each defined with two prices for revealed-identity and
- 11 two prices for anonymous identity. The pairs of prices are for accurate
- 12 information, and approximate information (with a minimal level of
- 13 approximation).
- 14 d) A mapping of data elements to price-functions.

15 Again, the price of a full query is computed as the sum price over all information  
16 requests to all data elements.

#### 17 **4.1.3 Superadditive Price Rules**

18 We can allow a non-linear interaction between the price of queries on individual data  
19 elements with the introduction of additive interaction prices, which are fired when a  
20 threshold is exceeded for the total amount of information released over all data elements  
21 accessed in the same query. This is a simple approximation to a more general  
22 combinatorial price rule, which would price all combinations of data access terms  
23 explicitly. For example, a rule could state: “If more than T data elements are accessed in  
24 set Sensitive then add a “bundle price” L to the total sum price of the query.”; or: “If the  
25 total weighted access to data elements in set Sensitive exceeds threshold T\_1, then add  
26 bundle price L\_1 to the total price of the query”; where the weighted-access is computed  
27 with weights equal to the degree of uncertainty associated with queries on elements.

#### 28 **4.2 Computing the Price of a Query**

29 A query can be quite complex, involving a number of operations on data stored in  
30 the data repository. Call a query session a sequence of operations performed in response  
31 to a query, and before the answer is returned to the agent that submits the query.





0969098-102700

1 separation is maintained between the data and the state of the program that executes  
2 queries.

3 Each query session is associated with a table of data record pointers. The table  
4 records valid data record pointers, that can be referenced in later query operations, and  
5 the current price of queries performed with the data record:

| 6  | Data Record Pointer |  | Current price |
|----|---------------------|--|---------------|
| 7  | -----               |  | -----         |
| 8  | P_1                 |  | \$0.10        |
| 9  | P_2                 |  | \$0.20        |
| 10 | ...                 |  | ...           |

11 The SDI query execution module performs another key book-keeping service, recording a  
12 list of data record pointers that have been used to compute a result in the program that  
13 executes a query.

14 For example, suppose that set PS contains a list of data record pointers, and a  
15 subsequent query operation makes the following command:

16 `_count = Count(PS, 'property X')`

17 which can be interpreted as “count the number of records in set PS with property X”. The  
18 query-execution module maintains a table of variable accounts, which records the data  
19 records that were used to compute a result:

| 20 | Variable Name |  | Associated data record pointers |
|----|---------------|--|---------------------------------|
| 21 | -----         |  | -----                           |
| 22 | _count        |  | P_1, P_4, P_6                   |
| 23 | _tmp          |  | P_3, P_8, P_10                  |
| 24 | ...           |  | ...                             |

25 Whenever the value of a variable is computed from the values of a number of existing  
26 variables, then the list of data record pointers is the union over the data record pointers  
27 which were used to compute the values of the other variables.

28 Eventually, when a result is reported at the end of the execution of a query, the  
29 price of the query is the sum of the requested payment for the data record pointers that  
30 were used to compute the information provided, as stored in the table of data record  
31 pointers.

Figure 3 illustrates the method diagrammatically. We maintain a “firewall” between data and the state of the query execution, with all access to the values in data records via data record pointers, and invoking price rules associated with data records. On the query-program side we maintain a table of data record pointers, and the cost of the cumulative information received from that data record under that pointer, and a table of variables and the data record pointers which were used to compute the values in the variables.

All manipulation with the values of variables on the “query-side” of the data firewall can be performed with no additional data-access charge, because data access charges are only incurred in computing values of variables, when that directly or indirectly requires that additional information be collected about the data records.

The only charges are incurred when the processor receives information. This information might be in the form of a number, that represents the count of records with a particular property, aggregate information in the form of a new collaborative filter, or raw data extracted from data records.

For example, SELECT operations are free until information in the data record is used to compute the value of something which is used in computing the final response returned in answering a query. A charge is accounted for against the data records which are queried and selected, but that charge is only levied if information is used.

#### **4.1.1 Internal Market within the Query Execution Module**

The data side of the firewall is also responsible for performing the cheapest queries possible. For example, if there is a request to retrieve 100 data records with property X, then the database retrieves the data as cheaply as possible based on the access-control rules of the owners of the information in the database.

All operations are implemented within an internal market. The internal market ensures that information is provided according to the price rules associated with data records, and also places providers of data in competition with each other.

The system of SDI implements an internal market, with queries executed as cheaply as possible. Both the SELECT and the PROCESS commands allow an associated restriction on size, e.g. “SELECT 100 records with property X”, or “PROCESS 10 of the records with data pointers in list X1 with method Y”.

1 In one greedy variation, each stage of a sequential query is executed as cheaply as  
2 possible. For example, for each operation the query is implemented by choosing the  
3 cheapest data records that satisfy the criteria. For example, the records that are selected  
4 are the 100 cheapest records to query that satisfy property X, and the 10 records in X1  
5 that are cheapest to process with method Y are selected.

6 This greedy method is not necessarily optimal for a sequence of queries, for  
7 example when the records that are cheapest for the current query are more expensive for  
8 a later query. Another variation allows a set of records to be maintained, with the  
9 cheapest records used in reporting the final result to the user. It might not be optimal to  
10 use the cheapest records at every point. For example, if a query is described as "SELECT  
11 100 records with property X", then "PROCESS with method Y", and "RETURN result",  
12 it can be cheaper to select more records (e.g. 1000) initially, process all the records, and  
13 then return the 100 records that are cheapest after the processing.

14 The method of storing the data record pointers in a table, along with the price  
15 required for information, and keeping a record of the data record pointers that are used to  
16 compute values during the execution of a query program provides a technique to compute  
17 the total price of a query.

18 In the previous section we described price rules, which are associated with data  
19 records and prescribe on a record-by-record basis the price of executing a query.  
20 However, only data records that are used in computing the final result that is returned to  
21 an agent at the end of a query receive payment.

#### 22 4.1.2 Query Optimization

23 Finally, consider the following two examples:

24 COUNT the data records with PROPERTY X  
25 and COUNT the data records with PROPERTY (not X)

26 In an optimal pricing model these queries both have the same cost. To focus ideas further,  
27 suppose that no data records have property X, such that without additional query  
28 optimization the cost of the first query is \$0, while the cost of the second query is the sum  
29 cost over all data records of reporting that they do have data property X.

30 To solve this problem we add a simple query optimization routine in the pricing  
31 module, which can test different formulations of properties (i.e. negatives), and the effect



1 In step 2 the data records that relate to profiles of users that have recently  
2 purchased a flight are now sorted in terms of their salary field. An algorithm to sort  
3 performs a sequence of comparisons between data records, with each comparison  
4 requiring that more information is revealed about the underlying salary of the profile, and  
5 changing the price associated with the data record. The prices are updated in the data  
6 pointer table according to the data record price rules, and the query program now has a  
7 list of ordered data record pointers.

8 Now, in step 3, the query program initializes a new set of data record pointers.  
9 This includes all the pointers for users that are willing to have their identities revealed.  
10 The pointers are maintained in order of salary. At this stage the entry for this new  
11 variable in the table of variable accounts only associates the data records with this new  
12 property.

13 Finally, in step 4, a new variable is initialized to contain the first 100 data record  
14 pointers from step 3. These are the data record pointers listed in the variable accounts  
15 table. These are the only data record pointers that receive payment as a result of this  
16 query, because the intermediate information provided by other data records remains  
17 within the SDI query execution module and is not released to the agent that performs the  
18 query.

#### 19 4.3 Discussion

20 The pricing model as outlined above has the following attractive property:

21 The price of a query is independent of the order in which operations are  
22 performed, and optimization of order is not necessary for reasons of price.  
23 This is nice, because it means that agents do not have to perform lengthy optimization to  
24 structure queries in cost-efficient ways. The pricing module also ensures that:

25 Agents only pay for information received as the result of executing a query.  
26 Finally, we implement an internal market, so that with all other things equal, data mining  
27 and other types of queries are executed as cheaply as possible:

28 An internal market favors data records with low cost pricing instead of records  
29 with high cost pricing, with all other things equal.

30 As an example, consider this query:

31 COUNT (SELECT 'FIELD1 = X' and 'FIELD 2 = Y')



The pricing module is a core component of the central SDI data warehouse. It has two functions: (a) to *predict the price of a query without executing a complete query*; (b) to aid in the allocation of payment given a price quote.

We build into the pricing module a method to allow an agent to make a quality/cost tradeoff before executing a query, and even to specify a restricted query that will respect an agent’s budget. Already embedded in the query-execution module is an internal market for data records, to ensure that cheaper data is used in preference to more expensive data access, given two identical data records. Furthermore, already embedded in the query-execution module is that agents only pay for the *provision of information*. Finally, the method computes a price for a query which is independent of the order of operations.

We propose the following method for providing initial price quotes to agents in SDI:

1. Compute a range of price quotes in response to a query, for different amounts of computation (which can be varied by limiting the input set of data records), and for corresponding values of a “quality metric” provided by an agent that submits a query.
2. Allow the querying agent to select an optimal price/quality/computation tradeoff, and then receive payment, execute the query, and provide the response.
3. Allocate payments to accounts in the eBank that correspond with data records which are used to compute the value of the result of the query, with payments adjusted if necessary to account for any error in the price quote.

### 5.1 Computing a Price for a Query

We allow a small *fixed price*  $F$  for an initial price quote, to cover the cost of computation and prevent agents saturating the server. The price module computes a *price quote* based on a statistical technique to compute the price of the full query without executing the full query.

The basic idea is to perform the query on a number of records, e.g. 1%, 5%, and 10% of the total records that will be used for the full query, and then extrapolate to the full query size. We need to perform the complete query because on a subset of data records, instead of part of the query on all data records, because query programs can have different phases—and each phase might incur very different information costs. The



1 number of records that are used for the query can be restricted by making a random  
2 sample of the total domain of data records used for a full query.

### 3 5.2 Allowing a Price/Accuracy Tradeoff

4 In computing the initial price quote we also suggest a tradeoff for the user  
5 between “result quality” and cost. *Quality is subjective, and usually best measured by the*  
6 *agent that requests a query on the data. Therefore, we suggest that the querying agent*  
7 *should provide the metric, but limit this metric to a scalar value to prevent information*  
8 *leakage.* We allow an agent in formulating a query to specify a key metric that it will use  
9 to choose an appropriate size of query to execute.

10 The core technique that we use to control the accuracy of the result of a  
11 computation is to limit the number of total data records that a query runs over as input to  
12 the query. Take a random selection of all the data, and use that for the processing. At this  
13 initial step we limit the price and the amount of computation performed.

14 As an example, consider a query to collect pseudonyms for agents with useful  
15 properties. The metric of interest in this case is the average per-pseudonym cost, and this  
16 can be provided by the pricing module *without providing any useful information to the*  
17 *querying agent.* Given this information the agent can then decide whether to proceed with  
18 a query, and how much money to allocate.

19 Similarly, perhaps a querying agent cares about the quality of match between data  
20 records and data records with ideal properties. The agent can provide a method to  
21 *instrument* its query to compute this quality, and the price module can in pricing the  
22 query provide guidance about the tradeoff between running the query over different  
23 numbers of data records.

24 Another more involved method is to introduce randomization as the query is  
25 processed, because some owners of information may provide randomized information  
26 more cheaply than exact information.

27 That will vary depending on the number of data elements that are queried.

### 28 5.3 Making a Price Quote for a Budget

29 Consider the problem of an agent that wished to perform a query with a budget *B*.  
30 For example, I would like to spend \$1000 to compute a collaborative filtering model.  
31 There is a particularly simple way to allow this:









- 1       \* Pagers
- 2       \* Wrist watches (typically one-way, receive only)
- 3       \* ATM machines, point of sale kiosks
- 4       \* Ceiling-mounted video cameras
- 5       \* room-based motion and heat detectors
- 6       \* Internet-enabled automobile
- 7       \* Chip-enabled appliances
- 8       \* Electronic Schedule Books

## Smart Home Application

The concept of a “smart home” is not a new one: Even in the 1950’s futurists discussed the possibility of imbedding intelligent systems into a house, relieving the homeowners of such boring daily tasks as watering plants or adjusting the thermostat. Such a system could easily be implemented today using control computers running simple sets of IF-THEN rules, but might soon prove unreliable because of its inability to adapt to new situations. For example, a garden party could easily be ruined if inadvertently scheduled at the same time as a scheduled lawn-watering.

LEIA provides a much more sophisticated alternative. Given a house that incorporates an internal local area network (LAN) linking household appliances, sensors, communication devices, and home computers, LEIA could easily form the basis for a truly intelligent and adaptive home environment capable of handling daily chores, security, and information delivery.

Firstly, LEIA could easily subsume the tasks traditionally slated for use in a smart home; however, its ability to infer the location and needs of family members would greatly enhance the handling of these tasks. For example, the above-mentioned garden party would not be spoiled because a motion sensor would inform the watering system that a party is taking place on the lawn. Or, being set in a power-saving mode, LEIA could light only those rooms currently being used; as a family member moved around the



1  
2       iii) appliances

3       \*alarm clock

4       \*light switches

5       \*stove

6       \*dish washer/clothes washer

7       \*water taps

8  
9  
10       In addition, accessed via secure channels (perhaps using a personal agent as an  
11 intermediary), three major types of personal data are also useful for running the smart  
12 home:

13  
14       i) Explicit rules for routine tasks (E.g., desired schedule for lawn watering)

15       ii) Expressed preferences (E.g., a desire for morning news about international  
16 events)

17       iii) Historical database of previous interactions between the family and the house  
18 (This will  
19 allow LEIA to learn daily routines and family habits).

20  
21       There are two major classes of data that need to be served by LEIA: (1) control  
22 signals for household systems and appliances, and (2) information desired by household  
23 members.

24  
25       To a large extent, household systems can be appropriately controlled using a  
26 traditional rule-base, although a fuzzy rule-base would be inherently more flexible (IF  
27 <the bedroom is somewhat chilly> AND <it's almost time to wake up> THEN  
28 <substantially increase the bedroom's heat>). Many of these rules would be explicitly  
29 programmed by the family from the beginning; others could be learned using inductive  
30 methods.





complementarity) to the print media viewed by the user presently. Alternatively, the data captured by the miniature camera device relating to the particular pages of print media presently observed by the user may consist exclusively of an identifier code a portion of which is used by the system for purposes of identification of the code (or alternatively such identification may be determined by its physical coordinates on the page itself). The other portion of the code is a unique identifier of the particular page, which may be used to identify manually by a content expert whereby select print materials and particularly their associated electronic counterparts are either matched, designed and/or customized for use within their system context of the present convergent media application. It is apparent that such a system could also be usefully deployed to benefit advertisers, e.g., presenting advertisements electronically to correspond with particular printed content or matching printed ads with corresponding electronic counterparts.

This complimentary content, be it audio, video, text, or combinations thereof, including advertising conforming to user-defined conditions, is delivered through various in-home devices located in the proximity of the user.

### **Smart Office Application**

Large business centers can be very confusing places: people are often away from their desks, many different meetings are scheduled and rescheduled during a single day, people are called off to emergencies, and clients drop in for unexpected visits. The inability to coordinate people and information in such situations leads to greatly reduced efficiency.

LEIA offers a unified solution in the face of this chaos; by linking together people's schedule calendars, communications systems, contact information, and location coordinates, a system can be devised to (1) shepherd a visitor (even if unexpected) around the business center, (2) push vital contact information about the visitor to those in the company who will be visited, and (3) adaptively coordinate important meetings within the company.

**When a visitor comes to the center, he will:**

b) receive a hand-held computer that, while presenting the visitor with useful information, will emit a constant signal enabling LEIA to pinpoint his position at all times.

On the employee side, LEIA can acquire signals from:

a) active badges

b) telephones

c) beepers

d) interactions with terminals

All sensors are connected directly to a LAN internal to the company. Only basic security measures need to be enacted.

Registered visitors will likely be profiled in standard contact files containing information about their own company, position, professional interests, previous interactions with the company they're visiting, and basic personal details (e.g., photograph, wife's name).

Employees of the company using this LEIA-based system will have much more information available. In addition to the standard contact information, LEIA will have access to their phone and e-mail logs, daily calendar/schedule, work group and project details, position within the company hierarchy, and work-station notes (e.g., an unexpected “out to lunch” message could be posted on-screen by the employee before they run off).

Although some higher-level intelligence is needed (e.g., for matching interest profiles between employees and visitors), the office-based system could well be run with



Given that multiple employees have signed up for a meeting, or that a single employee has sent out a signal indicating the need for a meeting and its priority (standard or emergency), LEIA gets to work coordinating the details. Cell phone and beeper signals give employee locations outside the office, active badges signal employees already at the office. Calendar agents are contacted for availability; a standard meeting might depend on people having time slots free, whereas an emergency meeting would be given priority status, and could bump items already scheduled (given that they have been assigned low priority).

An optimal meeting time and place (or teleconnection for remote employees) is set up, and any pre-meeting information is forwarded to the appropriate workstations.

meeting, engaged in an important phone conversation

## Resolution Credentials

Section 2.2.1 of this patent discloses the use of resolution credentials for the facilitation of agent-agent interactions. Resolution credentials enrich the quality of such interactions because they certify the status claimed by individuals, allowing agents to maintain relationships of trust. On a more practical level, resolution credentials are of use for matching and introducing users to each other on the basis of common interests, for guarding access to high-level users (with the agent acting as a gate-keeper), and for flagging times at which an individual is reachable (and not tied up in meetings, or the like, as specified by the scheduling agent).

## 7.2 General Description

We allow a user of SDI to push some data **A** to the data warehouse, and make this data available to all other agents via the SDI query-execution module, and to retain local control over other data **B** on the client machine. In this section we discuss this and other variations, and explain how the functionality of the central SDI data warehouse can be





In Figure 5 we illustrate the situation where agent 1 has its personal information about the user stored in the central SDI data warehouse, and provides agent 2 with the ability to process a query on the data warehouse and access information about the agent. The query might either request information about the agent’s profile, or provide a method to execute in the data warehouse, where the result of the method is an optimal action for Agent 2 to make in providing personalized information and/or services to the user with agent 1.

### 7.3.2 Data on the Client Machine (client-side SDI database) and in SDI data warehouse

Figure 6 shows an alternative variation where data is stored on an agent's client machine and in the SDI data warehouse, and the client-side SDI data warehouse can respond to queries from agent 2 in the same way as the central SDI data warehouse. In particular, as we suggest later in this document, this variation allows an agent representing an on-line consumer to store very personal information on a local machine, and for example only allow other agents to benefit from the *results* of analysis on that information without receiving the actual information.

### 7.3.3 All information stored Client-side.

Another variation on the basic SDI architecture, is that all information is stored only as *virtual links* within the SDI data warehouse, with physical storage on an agent's local client machine. Preferably, in order to leverage the value of information the client machine would be always networked with the central SDI data warehouse. A current technical solution which seems appropriate would use a *Set Top Box*, connected via cable TV lines to the SDI network. Many households in the US have set top boxes, and this



1 would provide a quite convenient way to leverage the value of information but keep it  
2 secure against inappropriate access in a central data warehouse.

3  
4 In this variation, to support a personalized session with another agent the data is stored  
5 locally, and can be processed using a client-side SDI data warehouse with a local query-  
6 execution module. To also support data mining by agents of a shared database, the client-  
7 side data records could useful push *data types* and *price rules* to the virtual shared  
8 database, to allow more efficient searching. This variation describes a *distributed SDI*  
9 *data warehouse*, with indexing performed over on-line client machines.

## 11 8. Distributed Query Execution

12  
13 In addition to allowing distributed data location, the system of SDI allows *distributed*  
14 *query execution*, via client-side query-execution modules. The advantage with this  
15 variation is that data can be physically stored on trusted local machines close to a user of  
16 SDI, and the queries can be physically executed on the machines so that not even  
17 intermediate results are available outside of a local firewall protected system. There is no  
18 technical solution provided within SDI to prevent agents sharing information to other  
19 agents, other than to prevent agents from receiving that information in the first place. We  
20 do this using three many techniques: pseudonyms, information randomization (see  
21 Section 10), and query-execution in the SDI central server without release of raw data  
22 inputs.

23  
24 We can implement safe client-side query execution with direct replication of the key  
25 functionality of the query-execution and price-module of the central SDI data warehouse  
26 on a user's client machine. A querying agent can push the same method to an agent's  
27 client machine as it would submit to the central SDI query execution module, and an  
28 agent can execute the method locally and release information consistent with his/her data  
29 release rules. The SDI client-side agent can be configured, via the profile management  
30 and data-release policies, to control the type and amount of profile information released

The challenge is to be careful that the results of a query do not reveal too much information, but this is solved using the same *price rules* as in the central SDI data warehouse. The mode of interaction supported is QUERY(PID), i.e. execute a query on a particular pseudonym ID, and the client side query execution module ensures that the vendor has a certificate from the client-side SDI agent to query the information in the local database about that pseudonym. The control over the profiles that a vendor can access, coupled with the price rules to ensure that information is perturbed as necessary and that not too much information is released, and to allow an agent to leverage the value of information.

A user's agent can retain control over local information by limiting the information that can be released in response to a query. For example, although a query can be complex and access a lot of profile information on a user's client machine, we suggest an additional protection that controls the information that a vendor can receive for future use. It is possible to limit the response that the vendor receives to a fixed message size, e.g. to 10 bits for example, to provide a very strong overall control on the ability of a vendor to use information again in the future.

### 8.1 Example: Customized Vendor Web Pages

The method, called *safe client-side query execution*, has a direct application to systems for personalized information delivery, where it is not desirable to release profile information to an information provider, but preferable to allow the information provider to provide methods (e.g. queries) that are executed on a client machine, and personalize information before it is provided to the end user.

A key application of this technique is in B2C on-line e-commerce where a vendor can push methods for personalization of a virtual shop front to a user, and the user's SDI enabled client machine can implement the methods (which are queries) on the client, and

1 push the result back to the vendor's server. The server then constructs and pushes an  
2 optimized virtual shop front to be displayed on a user's local display.

3  
4 At present advertising networks such as DoubleClick ([www.doubleclick.com](http://www.doubleclick.com)) are able to  
5 track a user as he/she browses across multiple sites in the DoubleClick network because  
6 the ad server can place a "cookie" text string on the hard drive of a user's client computer  
7 which identifies that user as he/she browses. SDI is designed to be used on-line in  
8 combination with a method to block cookies, replacing them for example with "safe  
9 cookies" [Netscape 96] to still permit stateful interactions during a single session with a  
10 vendor, for example "shopping basket" style interfaces.

11  
12 The SDI-client side proxy agent can replace the role that cookies take in automatic user  
13 log-in and password checking, through a Janus/LPWA-style implementation [BGGMM  
14 97; BGGMM98; GGMM98]. A user's log-in user name and password can be derived  
15 from a user's pseudonymous identity, and computed within a cryptographic framework.  
16 Janus also provides a technique to allow a user to receive e-mail pseudonymously.

17  
18 A core technique in SDI is to use a unique pseudonym for each vendor, and then  
19 selectively provide vendors with profile information from across multiple pseudonyms;  
20 either anonymously within the central SDI data warehouse for data mining purposes, or  
21 alternatively via client-side personalization so that a vendor can leverage a user's wide  
22 profile data without receiving *direct* access to that information.

23  
24 In providing profile information to a vendor during an interaction, so that the vendor can  
25 provide targeted products and servers to the user based on data mining that it has  
26 performed on aggregated user data in the central data warehouse, the user's agent must be  
27 careful to protect the identity of a user. This means that the agent must not reveal  
28 information under pseudonym P\_1 and information under pseudonym P\_2 to another  
29 vendor that allows the vendors to link the pseudonyms and reason that the agent  
30 represents the same user.

1 We describe in Section 10 the technique of random noise perturbation to release  
2 information in response to queries, and another simple technique is to carefully protect  
3 particularly sensitive information (e.g. social security numbers, etc.)  
4

## 5 8.2 Client-side Query Execution 6

7 As described earlier, a key variation of SDI retains a local data warehouse and query  
8 execution module that contains information specific to a single agent, on the client  
9 machine of that agent. It is then possible, via the same query execution controls as  
10 described in the central SDI data warehouse, to allow vendors to characterize and profile  
11 a user based on its local data, but without gaining explicit information about the user's  
12 local data. The vendor can push appropriate methods and targeted services. We describe a  
13 number of bottom-level applications later in the patent, including personalized web  
14 pages, ad-networks, etc.  
15

16 We describe one variation in which the vendor provides generic information to the client-  
17 side device, that filters that information locally. This is relevant for example in high  
18 bandwidth information services such as digital television and satellite systems, where a  
19 user's set top box can store program information locally and make local decisions about  
20 what programs a user might like to watch and when, without providing any profile  
21 information to the head-end server.  
22

23 In this variation an agent can submit a query to be executed locally on a *client-side SDI*  
24 *data proxy*, that stores information submitted by the local agent. The applications of this  
25 technique are exciting, because it allows a user to maintain even more control over  
26 profile information but still benefit from personalized interactions with vendors. The  
27 client-side SDI data proxy just contains data for the local agent, and allows agents to  
28 submit 'personalization queries' of the type 'execute this profiling rule and tell me what  
29 product to show the user'. The outcome of the processing might be personalized  
30 information, for example a selection of books that a user might be particularly interested

1 in based on its profile information and the methods passed to the agent by the agent that  
2 sends general information about the books in its catalogue.

3  
4 Figure 7 illustrates the process. The SDI proxy agent maintains a set of profiles for a user,  
5 relating to transactions that the agent has performed with other agents in the system, and  
6 also other information that relates to the user associated with the agent. A vendor agent  
7 can send generic information and a method for personalization, which is executed as a  
8 query on the local profile data base, with the results allowing a vendor to decide how to  
9 target its products and services to a user. For example, a book retailer might be able to  
10 provide a selection of books that a user might be particularly interested in based on its  
11 profile information and the methods passed to the agent by the agent that sends general  
12 information about the books in its catalogue.

13  
14  
15 We can provide *rule templates*, which are used to select the type of method that the  
16 provider of the information is requesting be applied to the generic information that is  
17 passed to the requestor agent. The provider agent only needs to provide *parameters* for  
18 the processing. For example, if collaborative filtering is the desired personalization  
19 technique then the providing agent provides a list of prototypical cluster centers, but  
20 does not need to specify a collaborative filtering algorithm. The interpreter takes the  
21 information and methods from the providing agent and selects the appropriate rule  
22 template to form a program that is executed as a *safe query* in the CPU, with the usual  
23 controls over the amount of data that a query can access.

24  
25 The output from this processing is personalized information, that can then be displayed to  
26 the user, perhaps to a final test in the control module for applicability.

### 27 28 **8.2.1 Preventing Information Leakage**

29  
30 The challenge is to be careful not to allow the *results* of a query, even those data is not  
31 explicitly revealed to an agent that submits a query, to carry sensitive information that a









## 9. Data Management

The invention of SDI, in its most general form, refers to the system for a privacy-protected market for data interchange between multiple parties. In this section we discuss possible methods for an agent to manage its disclosure of data to the data interchange, and to other agents in the system. We allow agents to adopt different identities, and submit different information under different identities. The idea is that a pseudonym prevents other agents from exchanging information. Given an identity management policy, the *data management policy* is then used to determine what data to submit to the central data server, or release to other agents, under a particular pseudonym. The agent needs to choose *what* data to release, and what *price rules* to associate with that data.

We propose client-side “proxy agents” to intermediate the interactions between agents, and manage data provided to the central SDI data warehouse; and methods to manage an agent’s identity in interactions with other agents, for example via pseudonymous and anonymous interactions. The proxy agents control the amount and types of information exchanged between agents. We refer to the proxy as “client-side”, because it resides on the machine local to the agents that participate in the system of SDI. The client-level SDI proxy, implemented as a client program running on the user’s client machine, manages all data release from the user to other agents and to the central SDI data warehouse. The proxy might usefully provide a rule-based interface to allow a user to select appropriate data management policies.

An important application of SDI is to a system for business-to-consumer (B2C) e-commerce, where SDI allows individuals to provide vendors with access to profile information that is collected client-side about an individual as he/she interacts with vendors, but retain control over the amount and level of detail that is made available and collect payment for that information. A key function of the client-level proxy is to control the ability of multiple vendors to track a user, by control of the identity of a user in its interactions and the information released. The client-side proxy agent maintains profile information about a user, but releases that information according to a user-defined policy.

1 The system of SDI, together with these client side proxy agents, allows a user to control  
2 the amount of information that is released to vendors and other agents, and allows users  
3 to exchange information within a privacy-protected and carefully controlled market place  
4 and receive personalized products and services The proxy might usefully provide a rule-  
5 based interface to allow a user to select appropriate data management policies. Vendors  
6 can run collaborative-filtering type applications on profile data from many online  
7 consumers, and use that information to personalize products and services in the future.  
8 Vendors can continue to build collaborative filtering models and other personalization  
9 models that used extended profiles, and users can receive the benefits of personalization  
10 without releasing profile information to a vendor; e.g. by personalization within the  
11 central SDI query-execution module or by personalization on a user's client machine.

## 12 13 9.1 Overview

14  
15 The basic assumption is that there is no technical method to prevent vendors from sharing  
16 information once they have that information, but it is possible to reduce the value of that  
17 information and the privacy implications for a user by making it very difficult to  
18 determine that two agents represent the same user in two different interactions.

19  
20 The ability to interact pseudonymously with other agents, and adopt a unique (temporary  
21 or persistent) pseudonym for each agent-agent interaction, is not in itself novel. This is  
22 described in the literature, for example as early as the work of Chaum [Chaum 85], and  
23 others more recently [FR 98; Cranor 96; BGGMM 97; BGGMM 98]. What is novel is the  
24 central SDI privacy-protected marketplace, that respects metainformation that agents  
25 associate with data, to carefully control the availability of data to other agents. Also novel  
26 is the method of 'distributed query execution' (Section 8), that allows an agent to push a  
27 method to another agent, and take a decision based on private information without  
28 learning the information.

29  
30 In Section 9.2 we present example identity-management policies. At one extreme, if an  
31 agent interacts with every other agent under a unique pseudonym, and never allows two

agents to match its identity based on information that it releases, then the other agents have no way of combining information about the agent. In the business-to-consumer e-commerce example, vendor A cannot share information about consumer Z with vendor B unless vendor B has some way of linking its customer with the customer of vendor A. This can be done at present via cookie mechanisms and advertising networks (e.g. DoubleClick), or simply by matching identifying information (e.g. user name, e-mail address, credit card number) across sites.

In Section 9.3 we describe data-release policies, which are methods that determine what information is released for each pseudonym a user selects, and the rules associated with accessing that information. We also note that an agent can choose to introduce a level of noise-perturbation to data before its release, so that the data is still useful but cannot be used to link an agent across multiple pseudonyms via the data that it releases.

Alternatively, an agent can release data with a reduced level of accuracy, for example within a range of values, and/or restrict the types of data that are released. Methods of data perturbation are discussed in more detail in Section 10.

Finally, note that there are two parts to the data management policy. One relates to the data that a user releases to the central SDI data warehouse under a particular pseudonym, and another relates to the data that an agent releases to another agent in a direct agent-agent interaction. The following bullets summarize the distinction, and note the sections for more details.

- [9.3: Data Management Policy] Manage the Release of Data to the Central SDI data warehouse (with either physical uploading of data, or release of links to physically distributed data); this includes the association of price-rules with data that is provided to SDI.
- [9.4: Agent-Agent Data Management] Manage direct Interactions with other Agents, for example controlling the identities used in interactions, and the amount



1 system [BGGMM 97; BGGMM98] allows a client-side proxy to compute a new pseudonym  
2 for an agent, and also a new e-mail address and user-name and password for that  
3 pseudonym, to allow access to web-pages that require user log-in. In particular, it is possible  
4 to compute new pseudonyms for a user without a central register of pseudonyms that could  
5 compromise a user's identity, through the technique of "blinded signatures" [Chaum 85].

6  
7 The identity-management policy can be configured by a user when he/she first registers with  
8 the system of SDI; for example a simple policy might classify agents according to the  
9 certificates that they can present, and select a pseudonym according to the agent's class.

10 Other reasonable identity-management policies include:

- 11  
12 • Absolute Privacy. At the highest level of privacy an agent interacts anonymously  
13 with every other agent so that agents cannot learn about the agent across sessions,  
14 and cannot personalize future interactions. Furthermore, information about the  
15 transaction with the agent, cannot violate a user's privacy when provided to other  
16 agents, so long as the agent does not release any identifying information during its  
17 transaction.
- 18  
19 • High Privacy. At the next level of privacy, an agent interacts with every agent under  
20 a unique persistent pseudonymous identity. This prevents another agent exchanging  
21 information about the agent with other agents, unless the agent reveals other  
22 identifying information. However, this policy does allow agents to receive  
23 personalized information over multiple sessions with the same agent, directly  
24 without the vendor using information submitted to the central SDI data warehouse,  
25 or stored on a client-side single-user database. Another agent can personalize  
26 information, but only on the basis of previous transactions with the agent. In an  
27 Internet environment this mode of interaction allows a vendor to track its customer  
28 across multiple sessions.
- 29  
30 • Medium Privacy. Use the same persistent pseudonymous profile with groups of  
31 other agents, perhaps segmented according to the type of task that an agent is

- Low Privacy. No attempt is made to protect the identity of an agent. This is the default privacy level provided via current Internet browsers, at least in the common usage of individuals. The system of Secure Data Interchange is useful with this level of privacy only to the extent of leveraging the value of information that has not been released to any agent. There is not technical method to prevent agents from exchanging information that relates to transactions they have performed with the same agent.

A general method to select a “medium” privacy strategy, with different pseudonyms for different agents, but some shared pseudonyms, is to suppose that agents present certificates. A certificate places an agent into a particular class, and each class is associated with a particular pseudonym. Therefore agents that can position themselves in a particular class that is shared with other agents can share profile information; while an agent that can not position itself within a shared class will receive its own unique pseudonym.

A useful default policy is to suggest that an agent adopts a unique pseudonymous identifier for each vendor that they interaction, and then explicitly link pseudonyms as they decide, via a link-command to the central SDI database. The policy is a hybrid of the core suggestions above. Initially a user selects high privacy, but the user can choose to incrementally relax this privacy decision on the basis of continued interactions with agents. A user might also



1 can be released, the types of queries that can be performed on that data, the price that  
2 must be paid to perform the queries. For example, we might hard code different price and  
3 data-release policy codes, e.g. A B C D..., and allow third-parties to provide maps  
4 between data types and an appropriate policy code.

5  
6 We view an agent's complete data release policy as an intersection of its data  
7 management and identity management policies. As with identity management, a simple  
8 variation will define a fixed number of data management policies, stating types of data  
9 that can be released and conditions, and a mapping from agent certificates to policies. An  
10 agent can receive whatever policy it presents a certificate for.

### 11 12 **9.3.1 Data Validation**

13  
14 The SDI client-side proxy can also provide a useful service of validating data in profiles.  
15 This can be important in scenarios in which it is possible that a user can try to misstate  
16 information in the database for personal gain. For example, consider a consumer-business  
17 e-commerce system in which vendors determine user discounts for products based on  
18 their profile information. Clearly all users will attempt to adopt the profile for the  
19 cheapest price, if that is known, and if that is possible. We can allow vendors to specify  
20 that they will only follow validated information, for example if a user's profile states that  
21 he/she is very price sensitive, then it is more useful to know that if the assessment is  
22 based on validated transactions performed via the SDI system. In general terms,  
23 information is valuable within a marketplace if the information is accurate, and we can  
24 use client-side validation to achieve that goal.

25  
26 Consider another business-consumer e-commerce example. In general it is useful to allow  
27 an agent to duplicate information across different identities in the SDI data warehouse,  
28 e.g. my Zip code might well be information that I am very happy to associate with all of  
29 my profiles because it is not too revealing of my identity. However, as a vendor I might  
30 like to know that when I execute a query I am not paying to receive duplicated data  
31 records for the same user. One useful way to prevent this is to allow a user to also state



1 with some records “this is the only user profile for which I am choosing to submit this  
2 piece of information”. For example, if user A purchases book X then the user might  
3 choose to associate information about that book purchase with just one of his/her profiles  
4 in the SDI data warehouse, e.g. the profile for the “book-reader”. Stating this allows a  
5 vendor to specify within its query that it is only interested in receiving information  
6 associated with profiles that include this “exclusivity” claim. The SDI client-side proxy  
7 can provide a guarantee that the user only associates data with one profile, for example  
8 using a cryptographic method, e.g. signing the data record with a key to indicate that it is  
9 a unique record.

### 11 9.3.2 Data Perturbation

13 In addition to an agent’s identity management and data-release policy an agent must be  
14 careful that it does not compromise its identity-management policy by releasing identifying  
15 information. For example, an agent must be careful not to release the same piece of  
16 information X under pseudonyms P1 and P2 if it is unlikely that the information would  
17 relate to two different agents. In the next section we discuss the concept of data perturbation  
18 in some detail. The idea is that an agent must add enough noise to any information that is  
19 released to prevent identification by another agent, unless the other agent already knows the  
20 agent’s true identity. This date-perturbation module can run on top of identity and data-  
21 release methods. Without careful management of the control of information a  
22 pseudonymous identity management policy is redundant. Remember, whenever  
23 pseudonymity is compromised there is no technical method to prevent vendors and other  
24 agents from exchanging information about the agent.

26 Agents are careful to reveal only information that will not allow a vendor to link the identity  
27 of a user across multiple pseudonyms, defeating the identity management policy.  
28 Pseudonymity can be broken whenever an agent reveals the same piece of information, X, to  
29 multiple agents, e.g. A and B, and that information X has significant discriminative power.  
30 Agents A and B might be able to deduce with quite high probability that it is likely that the

two pseudonyms refer to the same agent; and therefore combine their profile information and defeat the user's profile management policy.

Consider an extreme example in business-to-consumer e-commerce: if agent 1 reveals its social security number to vendors 1 and 3, its pseudonymity is broken. Vendors 1 and 3 can now communicate and deduce that the agent that interacts under pseudonym P1 and pseudonym P2 is the same agent, and therefore share profile information that the agent would like to isolate from each agent. The pseudonymity is broken in this case because social security numbers are unique identifiers. A slightly more subtle example is to consider an agent that provides a combination of profile information to two vendors, where the information taken together is suitably revealing that the vendors can conclude that the pseudonyms relate to the same user with high probability.

At the other extreme, we might simply provide no personal identifying information, such that in a business-consumer e-commerce application the vendor knows only the type of product that the agent is looking for, or the request for information made by the agent. In this situation we can use the methods disclosed in [SECTION XXX, END OF BOTTOM-LEVEL] to complete a transaction with pseudonymous physical mail and pseudonymous payments.

A slightly more advanced method is to classify information as identifying and non-identifying, where information is non-identifying if no amount of the information can break a user's identity; e.g. information common to many other users, perhaps the city in which a user lives, a rough salary range for a user, a user's banking institution, etc. In comparison, identifying information might include a user's street address, a user's favorite book, a user's exact salary, etc.

The data-perturbation approach is to select the information to release, via an agent's data-release policy, and then add a necessary amount of random noise to that information to protect a user's identity. We can add a small amount of random noise to data, enough to prevent identification, but not so much that the data has no value to another agent. For

example, my year of birth and Zip code might be almost as valuable for customization purposes as is my full date of birth and street address.

### 9.3.2 Click stream data.

One type of data that is especially interesting in a browsing environment is click stream data, which is stored at the client machine and represents a sequence of clicks that a user has executed, possible across multiple vendors. The data can be stored on a user's local client machine and periodically released under an appropriate pseudonym to the central SDI data warehouse. The client-level proxy server that runs on a user's host machine is in a unique position of being able to monitor the user across different pseudonyms and across different vendors' sites. The client proxy might also collect information about:

- The information that is displayed to a user (e.g. the text, the pictures, etc.)
- Information typed at the keyboard, and profile information transferred from the client machine to a vendor.

The data is gathered by passively observing the actions of the user, and not by direct question-and-response. Possible click stream data release policies include the following data-release policies:

- A. Release no information.
- B. Only release data on the URLs of the most recent sites visited.
- C. Release data about the URLs of the most recent sites visited, and the information displayed to the user.
- D. Release data about the URLs, the information displayed, and the information entered by the user.

in combination with a suitable identity-management policy. For example, an agent might state that all click stream data should be released under the pseudonym that the agent adopts for a particular transaction, i.e. all click stream data with amazon.com should only

1 be associated with my “buying books” pseudonym. An alternative policy might state that  
2 click stream data should be stored under a unique pseudonym, and not identified with any  
3 of a user’s other pseudonyms. The agent might then allow a vendor to perform limited  
4 queries on that data, for example only query information related to particular domain  
5 names. Many variations are possible: e.g. use a unique pseudonym for each new URL  
6 domain, i.e. whenever a user skips to a new site, submit click stream data under a new  
7 pseudonym.

8  
9 Click stream data can be subject to random perturbation, just like standard data-- for  
10 example removing time-stamp information and adding noise to the URLs that a user clicks.

11  
12 There may be other click stream data release policies worth considering outside the  
13 preferred embodiment discussed above. For example, Intermind’s patent number XXXX,  
14 entitled XXXX, provides for the release of a user’s information according to the terms  
15 and conditions of the user’s own data disclosure policy.

16  
17 Such an arrangement could be further enhanced by conditioning the quality of the data  
18 released by the vendor to the user on the strictness of the user’s own data disclosure  
19 policy -- users could then be rewarded to share more of their clickthrough data. Various  
20 levels of security could also be guaranteed to the user releasing such information – for  
21 example the data perturbation technique could be used to ensure that the vendor only  
22 receives aggregate information about his visitors.

#### 23 24 9.4 Agent-Agent Data Management

25  
26 It is important that an agent that submits data to the central SDI data warehouse also  
27 maintains careful control over the data that is directly released to other agents through  
28 one-to-one interactions. It is necessary to prevent a “black-market” in data. Furthermore,  
29 a user might simply prefer that another agent (for example representing a vendor) does  
30 not know certain pieces of information.

1 With respect to preventing a black market in information about a user, it is important to  
 2 prevent two different agents linking the identity of an agent under two different  
 3 pseudonyms, based on comparing information provided by an agent under each  
 4 pseudonym. We do not propose a technical solution to prevent these agents exchanging  
 5 information (outside of SDI) about the agent when this type of linking is possible. Such  
 6 violations can not only affect a user's privacy, but might also decrease the value of  
 7 information provided to the SDI data warehouse, since information can be readily  
 8 exchanged between agents outside of the constraints of SDI.

9  
 10 Fundamentally, the policy under which an agent handles data release to another agent is  
 11 no different from the policy with which data is released to the central SDI data exchange.  
 12 This is a special case, in which the agent knows that it is providing explicit permission  
 13 for an agent (specifically the agent with which it interacts) to receive all information that  
 14 it releases. As indicate above, the pseudonym and the data release to another agent can be  
 15 determined by certificates that the agent is able to present, and an agent's local data and  
 16 identity management policies.

17  
 18 For example, suppose that the policy defines that certificate A confers the right to receive  
 19 a persistent pseudonym, and also receive links to the pseudonyms in set P1. Suppose the  
 20 certificate B confers the right to receive a anonymous pseudonym, and receive no links,  
 21 and certificate C confers the right to receive a persistent pseudonym, and links to  
 22 pseudonyms in set P2. Finally, suppose an anonymous profile is the default. Now,  
 23 suppose a vendor presents certificates A and C. This vendor receives a persistent  
 24 pseudonym, and links to pseudonyms in the union of sets P1 and P2; a vendor that  
 25 presents A and B receives a persistent pseudonym and links to pseudonyms in set P1; and  
 26 a vendor that presents no certificates receives an anonymous profile and no links to other  
 27 pseudonyms.

28  
 29 9.5 Light Clients: Web-Centric Data Management  
 30





1 transaction via his/her credit or debit cards, for example on ATM machines and POS  
2 kiosks, providing promotional offers and coupons.

### 3 4 9.6 Smart Cookies

5  
6 Current practice in web-browser based consumer-to-business electronic commerce is to  
7 use cookies, which are identifiers placed on a user's hard drive, to identify a user across  
8 an extended period of time. For example, if I access the New York Times from my home  
9 personal computer on Monday, and then again on Tuesday, the New York Times server  
10 can identify that I am the same individual and build a profile of my interests, i.e. the new  
11 stories which I choose to receive first each day. This type of information about all of the  
12 users that read the Times can allow collaborative-filtering type techniques and  
13 personalization of information in the future, such that my "front page" is different from  
14 the front page of someone else.

15  
16 However, cookies have the unfortunate side effect of allowing an individual to be tracked  
17 across the web pages of different vendors, for example across the web page of the New  
18 York Times ([www.newyorktimes.com](http://www.newyorktimes.com)) and Amazon ([www.amazon.com](http://www.amazon.com)), if the Times  
19 and Amazon both embed content from the same third party in their pages. This happens,  
20 for example, with the DoubleClick advertising network. DoubleClick  
21 ([www.doubleclick.com](http://www.doubleclick.com)) operate a virtual network of pages, and can track a user across  
22 any page within their network, and gather a very comprehensive user profile. Although  
23 the cookie mechanism is designed so that only vendors with the same domain name can  
24 access cookies on a user's hard drive, they can easily be used to profile users across  
25 multiple vendors, for example with a double-click style network that embeds a universal  
26 advert server within each page.

27  
28 The system of SDI allows controlled personalization, such that a proxy-automated log-in  
29 session where the proxy presents a user's pseudonym to a vendor allows that vendor to  
30 track a user over time at its own web site, but a vendor cannot track a user across web  
31 sites. The SDI client-side proxy agent will disable cookies in their current form (although





BGGMM98;GGMM98] or W3C proposals [CR98; RC99; W3C-OPS 97]; essentially the pseudonym and an associated password that a user adopts for a vendor allows log-in to that vendor.

Also important when supporting a system that allows a user to browse pseudonymously is that the physical attributes of a network system are removed of their identifying characteristics. For example, another role of client-side data management agents is to strip the ‘from’ field in a HTTP/TCP message.

## 9.7 Implementation Details

In an Internet browser environment the client-side proxy agent that provides distributed data management for an agent might be implemented as a plug-in into the browser, that can for example be downloaded from a central SDI server. The browser is then configured to use the SDI proxy as its proxy, and the SDI proxy itself connects through a user's ISP (or other intranet gateway) to the Internet, and on to other vendors.

The user provides his/her SDI proxy with personal information, such as his/her name, mailing address, and e-mail address. The client-level proxy registers then registers the user with the central SDI server, providing the server with the name, address and e-mail address of the user. Other basic user information might include demographic information, for example a users job, marital status etc. The client proceeds to automatically generate a unique SDI user ID code, and a private key to allow future authentication of its log-in.

The client can create a unique public key/private key pair. This key pair can be generated only once for a person, and although the central SDI user ID server does not know the key pair, the server can verify that a key pair is only generated once-- because a new user must present proof of identity to establish an account. The client generates a unique user identifier, UUID, for example with the methods taught in [Chaum 85; Schneier 92]. The UUID can then be blinded and signed to certify that a user is registered with SDI, using Chaum's technique of blinded signatures so that the certifying agent does not the identity of



the aid of a decision tree an appropriate profile management policy for a user. Profile management policies define how a user will interact with various classes of vendors (depending on the nature of the business that the vendor is engaged in), the kinds of uses to which the transactional information that a vendor collects can be put to, and the amount of information that a vendor which collects profile information about a user is authorized to release. The client-level proxy manages a user's interactions with vendors, to keep them within desired policies.

### 9.7.2 Generating a New Pseudonym

The method of blinded signatures [Chaum 85; Chaum 92] provides a useful technique to generate new pseudonyms for a user, without any centralized database that stores pseudonyms. The client-level SDI proxy can generate a new identifier, comprised of a sequence of bits, that will be unique with a high degree of probability. The identifier is "blinded" and then submitted to a trusted-third party to be authenticated for use as a pseudonymous identifier. When authenticated the blinding factor is removed, and the final signed pseudonymous ID can be used as a new identifier. The signature can allow a user to associate certificates with that identifier, and also to verify that the user has only a single pseudonymous ID with a particular vendor.

Pseudonym administering authorities (PAS) cannot build dossiers of the pseudonyms, because users submit "blinded" identifiers. The only information that a PAS has is the list of unique vendors that a particular user has registered with. The pseudonym administering server can be operated by an agent with a trusted relationship with a particular vendor.

Every SDI user has a unique identifier, a UUID, that is presented to the PAS with the new pseudonymous identifier to be validated. The PAS can verify that this is the first pseudonym for a particular vendor. Each Pseudonym administering server has a public key / private key pair (PKPAS, SKPAS) for each Vendor for which it validates new pseudonyms. A signed pseudonymous identifier, signed with the private key of PAS, verifies that the pseudonym is valid. Cryptographic techniques ensure that the signature cannot be falsified, and allow







The system as outlined above can be implemented within the current HyperText Transfer Protocol (HTTP), as a sequence of challenge/response pairs between clients and servers. The HTTP Post/Response mechanism allows clients and servers to exchange data, and this data can be an instance of an XML Document Type, within the body of a HTTP message. The HTTP protocol is the underlying mechanism, with SDI messages contained in the body of the HTTP Post and HTTP Response as XML documents.

### Example: A Possible XML Representation of a User profile

The World Wide Web Consortium (W3C) SGML working group developed XML (extensible markup language) to provide an open and extensible grammar for structured data [XML]. An XML document has an associated schema definition to enable an XML-enabled browser to validate the structure of XML data automatically. A Schema in XML is called a Document Type Definition (DTD), and defines the names of tags, their structure, and their content model. XML allows the DTD for an XML file to be identified through a Universal Resource Indicator [URI] in the header of the file (see below). XML also allows URIs for mobile code resources to be referenced, in order to enable a client to process embedded XML data. An XML document must be well formed, and in order to be well formed the tags must form a tree structure. In addition, the DTD allows the structure of an XML document (an instance) to be validated against a particular schema. Senders and receivers must only send valid SDI files. Each SDI message is a valid XML document.

We provide an example XML instance and part of a Document Type Definition for use within the system of SDI. We assume in this example that profile information is represented







23  
24  
25  
26  
27  
28  
29  
30  
31

25  
26  
27  
28  
29  
30  
31

29  
30  
31

30

31





1       A. **[Strong Protection]** Assume that the adversary knows that an agent **A\_1** with  
2       data records **P** submits data into a database; i.e. assume that the adversary  
3       knows for sure that one of the data records in the database relates to agent  
4       **A\_1**.

5  
6       Strong protection is only possible if there are enough data records from other  
7       agents to allow agent **A\_1** to add noise to its own data and make it sufficiently  
8       like that one of the other data records is the data of **A\_1**; in particular the other  
9       data records must also be submitted with random noise perturbation, or rounded,  
10      such that it is possible that other data records are submitted by agent **A\_1**.

11  
12      B. **[Weak Protection]** Assume that the adversary does not know that the agent  
13      **A\_1** with data records **P** has definitely submitted data into the database.

14  
15      Weak protection is easier to achieve. It is only necessary to add enough random  
16      perturbation to data to make the number of possible data records over the  
17      population of possible data records that are supported with the perturbed data  
18      record large enough to prevent agent identification.

19  
20      Randomized data is still useful within SDI for data mining and other applications, so long as  
21      the amount of noise which is added to records is small in comparison to the value of a  
22      record. For example, we can still perform correlation across fields with randomization so  
23      long as the randomization does not destroy any trends between fields. Randomized data is  
24      marked as such within SDI, and labeled with the degree of degradation, so that SDI can be  
25      aware of the number of records to get relevant accuracy levels, and can report accuracy to  
26      customers. There is a tradeoff between the level of privacy protection and the level of  
27      aggregation at which responses to queries become accurate. It is possible to add random  
28      noise to data but still allow data that is aggregated across multiple records to be quite  
29      accurate, so that useful data mining can be performed. Binning or rounding of data does not  
30      have the same effect. For example, assuming additive noise and additive aggregation, then  
31      randomized data from a number of agents can be aggregated to obtain an aggregate value









1 **Pr(true-value = X)** This is computed as the sum of  $\text{Pr}(A_i) \times \text{Pr}(X | A_i)$  over all data  
2 records, and represents the probability that any one of the data records was represents data  
3 with true value **X**.

4  
5 Agent **A\_1** can now select parameters for distribution **f** based on this analysis, so that the  
6 probability that it is the agent to submit the new data record is less than **eps**, its desired  
7 protection level.

8  
9 The rule has the right behavior—the more data present in the system then the more accuracy  
10 an agent can use to submit its own information, for the same privacy level epsilon. Notice  
11 that the agent is more protected as:

- 12
- 13 1. **Pr(A\_1)** decreases, i.e. with more data records in the database
  - 14 2. **Pr(true-value= X/ X', f)** decreases, i.e. with more noise perturbation **f**
  - 15 3. **Pr(true value = X)** increases, i.e. there are more data points from other agents that  
16 might have true value **X**.

17  
18 The parameters for **f**, for example the mean and standard deviation in the case of noise  
19 generated from a Normal distribution, can be selected to set  $\text{Pr}(X' | \text{true-value} = X, f) = \text{eps}$   
20 for the case that  $X' = X$ , i.e. when the random noise is zero. This presents the worst-case,  
21 assuming that distribution **f** places the most probability on zero noise.

22  
23 Finally, to compute parameters for **f** the agent requires aggregate information about the data  
24 in the database that has been submitted by other agents. In particular, the agent needs **Pr(X)**,  
25 the probability that any of the current data records could correspond to data with true value  
26 **X**. In some cases it is important to receive this information without revealing true value **X**.

27  
28 In a simple case, for example when a trusted intermediary such as SDI stores information  
29 provided by agents, the agent can simply poll SDI for **Pr(X)** before selecting a level of noise  
30 to use to perturb its information. Alternatively, when SDI is computing the level of noise  
31 perturbation to add dynamically as information is released in response to queries then the

1 system of SDI can compute  $\Pr(\mathbf{X})$  directly. In a more complex case, consider a problem  
2 where an agent is releasing information directly to an adversary, and the relevant set of data  
3 points are data that the adversary agent has already collected from other agents. In this case  
4 the provision of information  $\Pr(\mathbf{X})$  must be done within a secret-protocol where the  
5 adversary does not learn the value of  $\mathbf{X}$  in the process. A straightforward way to achieve this  
6 is for the “adversary” agent to post aggregated information about the probability that a data  
7 record in its population has true value  $\mathbf{X}$ , for a range of different values of  $\mathbf{X}$ , and allow an  
8 agent to use anonymous look-up in a table.

9  
10 In summary, the following procedure can be used to select an appropriate level of noise to  
11 add to a data point:

- 12  
13 a. Choose a distribution family (e.g. Normal, Uniform, etc.), and let  $\mathbf{Par}$  denote the  
14 parameters that define a specific distribution.
- 15 b. Choose a level of privacy protection,  $\epsilon$  where  $0 < \epsilon < 1$ .
- 16 c. Request  $\Pr(\mathbf{X})$  from the database, i.e. the current probability that a perturbed (or  
17 otherwise) data record in the database has true value  $\mathbf{X}$ .
- 18 d. Compute parameters for the distribution to set  $\epsilon = \max_{\mathbf{X}'} \{\Pr(\mathbf{A}_1 | \mathbf{X}', \mathbf{X}, \mathbf{f})\}$   
19 where  $\mathbf{X}'$  is the value generated from  $\mathbf{f}$ .

## 20 21 **10.1.2 Examples: Strong Protection**

### 22 23 (a) Uniform Additive Noise Perturbation Distributions

24  
25 Assume that every agent submits information from a uniform noise distribution, centered  
26 around its true value. The decision variable in choosing a level of data perturbation when  
27 submitting new information is the range of the uniform distribution.

28  
29 Suppose agent 1 submits data point  $\mathbf{g}_1 = [4, 6]$ , to denote that its perturbed value is  $\mathbf{X}'$ , and  
30 the value was computed with additive uniform noise  $\mathbf{U}(-1, 1)$ . Agent 2 submits data point

1  $g_2 = [4,6]$ , agent 3 submits  $g_3 = [3,5]$ , agent 4 submits  $g_4 = [3,6]$ , and agent 5 submits  
2  $g_5 = [2,4]$ .

3

4 Suppose that an adversary knows that user John has true value  $X = 2.5$ , and that an agent for  
5 John has submitted a data record.

6

7 In this case the adversary can be sure that agent 5 represents John, because  $\Pr(A_5 | x = 2.5)$   
8  $= \Pr(A_5) \Pr(x = 2.5 | [2,4]) / \Pr(x = 2.5) = 0.2 * (1/2) / (0.2 * 1/2) = 1$ .

9

10 In comparison, for an adversary that knows that Mary has true value  $X = 5$ , then  $\Pr(A_1 | x$   
11  $= 5) = 0.2 * (1/2) / (0.2 * 1/2 + 0.2 * 1/2 + 0.2 * 1/3) = 0.375$ , and similarly for agent  $A_2$ . The  
12 probability that agent  $A_4$  is Mary is 0.25.

13

14 Now, a new user, Bill, with agent  $A_6$  wants to provide information about its data  $X = 5$ ,  
15 and wants to be sure that an adversary cannot determine its identity with probability greater  
16 than  $\epsilon = 0.1$ . Let  $d_6$  denote the uncertainty selected by agent  $A_6$ , and compute an  
17 optimal  $d_6$  as so that  $\Pr(A_6 | x = 5)$ , i.e. the probability that agent  $A_6$  has true value  $x =$   
18 5. This is computed as  $\Pr(A_6 | x = 5) = 1/6 * (1/d_6) / (1/6 * 1/2 + 1/6 * 1/2 + 1/6 * 1/3 +$   
19  $1/6 * 1/d_6) = 1 / (4/3 * d_6 + 1)$ , and is less than  $\epsilon$  for  $d_6 > 6.75$ . Therefore agent  $A_6$   
20 should generate a perturbed value with additive noise computed with uncertainty 6.75, and  
21 can then be sure that its identity is protected.

22

23 (b) Normal Additive Noise Perturbation Distributions

24

25 Now, assume a normal noise distribution, and let  $sd_i$  denote the standard deviation selected  
26 by agent  $A_i$  for its noise distribution. The analysis is slightly more complicated, because  
27 the  $\Pr(X | A_i)$  is now a function of the position of  $X$  within the distribution, not just  
28 whether it is in range as with uniform distributions.

29

30 In this case an agent chooses the standard deviation for its Normal distribution by assuming  
31 that the randomized value  $X' = X$ , i.e. that the random noise distribution generates zero



1 number of intervals, according to the total size of interval specified and the size of the  
2 current interval.

3  
4 A useful variation on this method is to control the amount and type of information released,  
5 but only releases accurate information. The degree to which a user can be identified if it  
6 submits information **X** to a vendor depends on the other agents that have submitted similar  
7 information. For example, if **X** is very detailed (e.g. my street address), then it is probable  
8 that no other agent has submitted the same information, and I will self-identify myself with  
9 that information. On the other hand, if **X** is quite abstract, for example my ZIP code, then it  
10 is quite possible that a number of other agents with the same ZIP code have already  
11 submitted and released the same information. The method discussed above can be simply  
12 adapted to allow a client-side SDI proxy agent, or the SDI data warehouse, to decide how  
13 much information to release in order to protect a user's identity, based on information about  
14 the data already in a database. For example, instead of my street address my agent could  
15 choose to submit my ZIP code.

### 16 17 **10.2.1 Example: Binning Techniques**

18  
19 Suppose that an agent must decide whether or not to reveal its salary on the basis of  
20 information that has been revealed by other agents. The database supports intervals:  
21 **<\$20000, \$20000-40000, \$40000-100000, >\$100000**. The number of agents with each  
22 prototype value is currently **10, 12, 16, 5**.

23  
24 Now, suppose a new agent represents a user with salary **\$200,000** and privacy threshold **eps**  
25 **= 0.1**. The prototype values prevent an agent from stating the high value (**> \$100000**)  
26 because there are less than 10 agents in the bucket. However, if the provider also allows  
27 agents an option, for example "my value is in bucket 40,000 – 100,000 or > 100000", then  
28 the agent can state that information safely. In this case, the probability that the agent is an  
29 agent with salary \$200,000 is computed as its weight in the top valued bucket, divided by  
30 the total weight, i.e. **0.5 / 5.5 < 0.1**

1 This technique of binning information, or rounding information, is simpler to implement  
2 that random noise perturbation because probability distributions must neither be sent to the  
3 provider or maintained by the provider. However, the information is represented with a  
4 static accuracy which is hard to improve as the amount of data submitted increases. This  
5 occurs automatically in the standard version.

### 7 10.3 Adding Random Perturbations to Discrete Data

9 Adding noise to discrete values is a little more complicated, because the new value must  
10 remain feasible. For example, in randomizing the name of a CD the new name must be the  
11 name of another CD, not some "made up" name. We need to add noise to make data  
12 elements "close" to the accurate values. With discrete data, such as the name of an artist,  
13 "close" must be defined within the correct metric. The appropriate metric is such that a  
14 "close" value shares many of the same characteristics. For example, it is not appropriate to  
15 assign a close value on the basis of a shared last letter in the first name, but it is appropriate  
16 to assign a close value on the basis of an artist from the same genre of music --- from a  
17 "semantic cluster".

18 Agents use discrete probability distributions to randomize data points.

19  
20 For example, suppose that an agent wishes to reveal the name of the artist that recorded the  
21 last compact disk that a consumer purchased. An artist's name is best viewed as a point in  
22 "artist space", and therefore as a discrete value. It does not make sense to change a random  
23 set of letters, because the new "name" will not be the name of a valid artist. Instead, the  
24 concept of noise is to randomly choose a new name close to the current name, i.e. choose a  
25 new feasible location in artist space. One reasonable solution for the names of recording  
26 artists is to define a neighborhood of artists that are close to the original artist, and select a  
27 new artist from within the neighborhood with equal probability. The metric that defines how  
28 close artists must be computed using a system that is common knowledge to the provider of  
29 information and the consumer agent, because the provider of information must be able to  
30 compute the distribution over true artist name, given a randomized name and probability  
31 distribution. We can define the probability distribution with a threshold in the distance













1 integrated into an existing advertising service which they operate and maintain, in which  
2 case the vendor receives an appropriate fee for data which is exchanged between his/her  
3 existing customers, and a reduced fee (which may be split with the central SDI service)  
4 for data which is exchanged by/between a member of his/her SDI service and vendors  
5 who are members of the central SDI service but not of his/her local SDI service. Unless  
6 or until the client-level proxy server becomes a standard or a large critical mass of end-  
7 users adopts the service, one of several compelling business models could be used by a  
8 vendor to encourage the visitors to his/her site to adopt the user-centric SDI service based  
9 upon the monetary incentives the user may receive for subscribing. Typically the  
10 resulting revenues are split between the user, the user-centric SDI service and the vendor  
11 (in exchange for promoting the service to their site visitors a share of resulting revenues  
12 generated may be necessary).

#### 13 14 11.1 An open SDI system

15  
16 An ISP level proxy server can contain the user profile generation module, profile  
17 processing module, user profile interest summary generation module and target object  
18 generation module which operate in distributed manner. This enables an ISP to  
19 independently implement the core functionality of the system without the cooperation of  
20 information vendors (Web sites) or their operators (Web hosts) who opt-out of SDI. The  
21 modules in third-party SDI servers can share information with the modules in network  
22 vendor servers. This flexible architecture enables the user-centric SDI service to be  
23 implemented (by ISPs or completed independently) and when available inter-operating  
24 with the complete data sets available from the information vendors.

25  
26 SDI can allow third parties to operate their own secure advertising and/or electronic  
27 commerce-based product syndication affiliate network (for all customers). In accordance  
28 with the preferred implementation of SDI, these "advertisements" are represented at the  
29 item-level and may be distributed across the (proprietary or main) SDI system network  
30 presented in the form of ads, affiliate or portal links to purchasables or sites (which may  
31 contain target objects as purchasables) and/or simply (transparently integrated) selections









6 This architecture also may be useful and is ideally suited for cross vendor product  
7 advertising as through an ad network or product syndication network using affiliate links.  
8 In addition to the user profile generation module, a target object profile generation  
9 module should also reside across the network vendor servers such that it is possible to  
10 generate target object profiles for target objects on network vendor servers. Alternatively,  
11 user profiles and target object profiles are downloaded to the client level proxy which  
12 performs collaborative filtering tasks as the user browses from site to site.

In both of these cases, the main SDI server can receive user profile data generated from the user profile generation module located on the ISP-level proxy, and target object profiles generated from the target-object profile generation modules located on the various multiple information vendor servers.

21 12.1 Support for Pseudonymous electronic mail

The ISP-level proxy server is positioned just behind the firewall of the user's local dial-up network (ISP or Intranet). The proxy provides protection for users operating under pseudonyms from point-to-point attacks and HTTP header-tracking by stripping HTTP header-information and forwarding HTTP packets on to their destination with no information other than their source at the ISP-level proxy server. The ISP-level proxy also supports pseudonymous e-mail, between users, and between users and vendors.

1 Figure 2 shows a couple of users connected to clients, that are in turn connected to the  
2 Internet through a local intranet, such as the network of an Internet Service Provider (ISP).

3 The proxy "washes" outgoing messages of any information that would compromise a user's  
4 pseudonymity, for example the "referral" field that contains the previous URL of a user in a  
5 HTTP message. HTTP messages also leak other information, for example browser software  
6 on a user's client machine, the operating system and a user's IP address.

7 A user can receive electronic mail through the PID and associated IP address of the ISP-  
8 level proxy server.

9 The preferred implementation of this system allows the user to periodically check for new  
10 mail. The client-level proxy gains access to the mail box that is associated with a  
11 pseudonym by providing a correct response (signature) to an ISP-generated challenge.  
12 Notice that with this solution, the ISP-level proxy has no way to connect the pseudonyms of  
13 a user, so long as the user's client is not identified in its messages to the ISP-level proxy  
14 server other than by the PID that the proxy makes a request for.

15 We can extend this mechanism using a technique taught in the Lucent Personalized Web  
16 Assistant (LPWA). The LPWA [BGGMM 97; BGGMM98] provides for a sequential access  
17 mechanism to the mailboxes that belong to a user through a one-way function that takes the  
18 user's SDI log-in name and password, and an integer from 1 to N, and computes the  
19 mailbox location. The mail server does not need to maintain a list of pseudonyms for each  
20 user, because the user is able to efficiently access all of its mailboxes sequentially as a  
21 function of other information.

1 Another variation, that relies on the user placing trust in the ISP-level proxy server, provides  
2 the ISP-level proxy with the e-mail address for each pseudonym. This push method is more  
3 efficient, because the ISP proxy and the client proxy communicate only when new  
4 messages arrive, but provides the ISP proxy with information to compute all the  
5 pseudonyms for a single user—probably undesirable.

6

## 7 12.2 Support for Pseudonymous Physical Mail

8

### 9 12.2.1 Vendor to User

10

11 A vendor must hold a “physical mail certificate” to be able to send mail (packages,  
12 letters) to a user under a pseudonym. The certificate is similar to the “electronic mail  
13 certificate”, in that it is signed by the private key of the user’s pseudonym, and indicates  
14 that the vendor with public key  $P^*V$  can send mail to the user (under the pseudonym).

15 Each user has a trusted physical address authority, just as it has a trusted electronic mail  
16 authority (the second-level proxy server), that maintains the physical mailing address for  
17 each pseudonym. When a vendor has a letter  $X$  to mail to user with public key  $PKP$ , the  
18 vendor generates a unique ID for the package,  $IDX$ , and sends the ID code and the  
19 physical mail certificate to the trusted physical address authority of the user.

20 The physical address authority receives the certificate,  $S(PKP, PK^*V, SEND\_MAIL)$ ,  
21  $SKP$ ), that indicates that the vendor is authorized to send mail to the pseudonym, and the  
22 packages identify code, signed by the vendor to certify that the vendor holds the secret  
23 key that matches the public key in the physical mail certificate.

24 The vendor then passes the letter  $X$  and the signed ID code to a trusted mailer, that  
25 supports pseudonymous mailing, and has been certified by the central SDI server as such.  
26 The trusted mailer then provides the signed ID code to the physical address authority,  
27 signed with the private key of the trusted mailer. The physical address authority verifies  
28 that the trusted mailer is a valid service, and releases the real address of the user to the  
29 mailer. The mailer now has the letter  $X$  that the vendor wants to send to the user with  
30 pseudonym  $P$ , and the physical mailing address of the user - and the package can be  
31 mailed. At not time did the vendor determine the true mailing address of the user, unless  
32 it works in collusion with the trusted mailer, but the trusted mailer is certified by SDI,  
33 and also audited by the chosen physical address authority of the user. The address  
34 authority will only release addresses to reputable pseudonymous physical mail agents.

35 We can operate physical mailing lists in the same way, and gain additional security by  
36 never releasing the pseudonyms or the mailing addresses to the vendor that has requested

1 the targeted solicitations. We can use a technique that is similar to the technique that we  
2 used for virtual mailing lists. The vendor describes its solicitation to the central Secure  
3 Data Interchange, which leverages as much data as possible (without violating the  
4 privacy policies of any of the users or vendors that are represented within the data). The  
5 central SDI server generates a list of suitable pseudonyms, and then provides a series of  
6 unique codes to the vendor, that the vendor can supply to its chosen pseudonymous  
7 mailer with the material that is to be mailed. The central SDI server also provides the  
8 appropriate address authorities with authorization to release the physical mail addresses  
9 to the mailer when presented with the IDs. Notice that at no stage did the vendor have the  
10 pseudonyms or the mailing addresses. The parties all have only as much information as is  
11 necessary - the vendor needs some way to identify its packages to the pseudonymous  
12 mailer. The mailer needs an identifier to present to the address authority, and receives the  
13 addresses. The address authority just needs to know what addresses to release and to  
14 which third party.

### 15 **12.2.2 User to Vendor mail**

16  
17 The Secure Data Interchange system also provides a mechanism for users to send  
18 physical mail to vendors that are registered with SDI with pseudonymous return  
19 addresses. In particular, when a user sends mail to a vendor, the first-level proxy server  
20 provides a tool that: (1) Computes/Looks-up the appropriate pseudonym for the user with  
21 this vendor. (2) Generates a unique ID, and submits a signed message to the central SDI-  
22 server, where the message relates the pseudonym, the vendor, and the ID. (3) Provides  
23 the unique ID to the user.

24 The user writes the unique ID on the envelope, and mails it to the vendor. Should the  
25 vendor wish to reply to the user, then the vendor can take the envelope to a  
26 pseudonymous mailer, and request that the envelope be mailed appropriately. The  
27 pseudonymous mailer verifies the identity of the vendor, and then submits the ID,  
28 together with the vendor's signature, and its own signature, to the physical address  
29 authority that is maintained by SDI. SDI releases the address to the mailer that can then  
30 return the mail.

### 31 **12.3 Pseudonymous Payment Mechanisms**

32  
33 The Secure Data Interchange architecture must be able to support all the standard electronic  
34 commerce functions that we take for granted, but while maintaining pseudonymity for users  
35 and following privacy policies. There are various different solutions to this problem.

#### 37 **12.3.1 Anonymous Credit Card Payment [LMP 94]**

11 When the vendor submits its “right to payment” and proof of identity to the second-level  
12 proxy server the proxy server first runs the charge through the user’s credit card, and if  
13 that clears, runs the charge from the vendor through the account of SDI (which could also  
14 be a credit card, or could be operated as electronic cash or some other mechanism for  
15 payment).

16 This “anonymous credit card” payment method has the following properties:

- 17 1. The user's credit card pays \$x, but does not know who receives the money except  
18 that it is going to the Secure Data Interchange.
- 19 2. The vendor receives payment for \$x, but does not know the user's credit card  
20 information, or the user's identity.
- 21 3. The Secure Data Interchange incurs no financial risk because it receives payment  
22 from the user before making payment to the vendor, although there could still be  
23 problems if the user complains about the quality of the good for example.

25 This protocol is simpler than full cryptographic anonymous credit card mechanisms  
26 because the SDI acts as a trusted third party to both the user and the vendor.

### 28 12.3.2 Electronic Cash [Chaum 85; Chaum 92]

Electronic cash is anonymous, just like physical cash. The user purchases electronic cash from an electronic bank, presenting blinded notes, so that the bank has no record of the note numbers that it issues to the user. For example, the user generates a new note number,  $X$ , and has the bank sign a blinded copy with its \$10 signature,  $S(B(X), SK_{BANK\$10})$ . Then the user, or the first-level proxy for the user, removes the blinding factor, and can use the electronic cash as tender. Whenever the note changes hands the recipient needs to check with the bank that it has not yet been spent, because notes are easily copied, but not forged.

Electronic cash has the same useful properties as anonymous credit cards, although it is perhaps a little more exotic. In particular, notice that the bank does not know to whom, or for what, payment has been made. and the vendor does not know which user made the payment - it just receives the payment. We have minimized the amount of information exchange that takes place between the various parties in the system.

## 1 12.4 Client-Side SDI Proxy

2

3 The client-level SDI proxy, implemented as a client program running on the user's client  
4 machine, manages all data transfer between the client machine (and the user), and other  
5 vendors and the central SDI data warehouse. A key function of the client-level proxy is to  
6 implement profile management for a user, to control the ability of agents to track a user as  
7 he/she interacts with multiple vendors. The client-level proxy also controls release of profile  
8 information: the addition of demographic and other personal information to profiles, and the  
9 control of random perturbation to fields to prevent linking across user profiles.

10 The client-level proxy maintains profile information for a user's collection of pseudonyms,  
11 and allows the user to view and challenge profile information. The proxy also provides a  
12 rule-based interface to allow a user to select appropriate privacy/personalization policies.

13 The primary mechanism that protects the identity of a user across multiple vendors and  
14 service providers is the ability to interact pseudonymously with vendors. The user can  
15 choose a unique pseudonym for each third party with which he/she interacts, and be  
16 absolutely certain that he/she is the only party that knows his/her true identity. There is no  
17 way that a vendor can know anything about the transactions that a user has had with other  
18 vendors under alternate pseudonyms unless the user chooses to disclose the equivalence of  
19 pseudonyms, or use the same pseudonym across multiple vendors.

20 It is useful to distinguish three key modes of use of the Secure Data Interchange system:

21

- 22 • Static data-mining. Query execution with no dynamic requests for new  
23 information from information providing agents.
- 24 • Interactive data-mining. Query execution that includes dynamic attempts to  
25 request additional information from information providing agents.

- Client-side data-mining. Query execution that is performed on the local client machine of an agent, based on data stored exclusively on that machine.

All three modes may use distributed information, i.e. it is possible that the information is stored in the central SDI database, or on distributed client-side information servers, or in third-party servers. The first 'static' case and the second 'interactive' case are distinguished from the third 'client-side' case in that the query execution is performed centrally in the SDI data warehouse in the former, and on an agent's client machine itself in the latter. Client-side data mining has particular application to privacy-protected customization of information and services in on-line business-to-consumer applications.

## 1. Static Data Mining

In static data mining queries are executed on the information that is currently present in the SDI data warehouse, and there is no opportunity to contact agents and request more information. Applications of this type of static data mining include all types of "standard" database queries, where it is assumed that the data set is static. Queries may be open-ended, i.e. "find me all data records of this type, and perform the following operations.." or closed, i.e. "perform the following query on the data record for agent with pseudonym P1". Queries may also have side-effects, i.e. "find all records that satisfy this constraint, and then take action A".

## 2. Interactive Data Mining

In interactive data mining multiple agents may be contacted by the system of Secure Data Interchange in the process of executing a query, to request new information from agents, or push information to agents. A central application of this mode of data mining is *matchmaking*, which is a process where information flows between agents if and only if both the profiles of both agents are mutually compatible. In interactive data mining the querying agent does not need to know the identities of agents that are contacted by SDI, this is all transparent, and hidden from the querying agent.

### 3. Client-Side Data Mining

In client-side data mining the querying agent executes a query with the information associated with a specific agent, and the information remains located on that agent's client machine throughout the query. The query is executed by providing the *query method* to the client machine, processing the method with local information, and then returning a response or taking an appropriate action (e.g. displaying product X for price Y). A central application of client-side data mining is to *privacy-protected customization*, where a vendor wishes to customize its products and services for a particular agent that is registered with SDI, and take advantage of personal information relating to that agent that is not generally available. As another application, we describe a client-side advertising auction, where advertisers compete for the right to display a banner advert to a user, based on local information about the user's preferences.

#### 4. Applications: Very Brief Overview

In the next section of the SDI description we describe some specific variations and systems that can be implemented within the general architecture. In overview, we describe the following key applications:

- **Safe user profiling and personalization.** This allows on-line users to receive personalized information and services without providing personal information to vendors, so that users retain control over their personal information. Users can interact with vendors under different pseudonyms, and provide information to the central SDI data warehouse to allow data mining. Finally, users can allow specific vendors to execute queries, where the result of the query is information that allows that vendor to customize its service.
- **Client-side user profiling.** The client-side SDI proxy can monitor the browsing behavior of a user, and submit data periodically to the central SDI data warehouse with appropriate meta information to provide the user with guarantees about the type of information that can be released to vendors during query execution.

- **Client-side user profiling.** The client-side SDI proxy can monitor the browsing behavior of a user, and submit data periodically to the central SDI data warehouse with appropriate meta information to provide the user with guarantees about the type of information that can be released to vendors during query execution.



- 1       • Static Data mining applications. Agents can submit queries to the data warehouse  
2       query-execution module and perform data mining and collaborative filtering on  
3       aggregated and anonymous information provided by data submitted to the  
4       warehouse by agents.
- 5       • Interactive data mining applications. Agents can request actions from the SDI  
6       system if certain conditions are found to exist in information, for example SDI  
7       can send information to other agents about services or products, if good matches  
8       are found. Another example is a request that a certain number of agents with  
9       particular properties be contacted and asked to take a particular action, we give an  
10      example within a transportation domain.
- 11      • Pro-active data mining. The system of SDI might itself pro-actively execute data  
12      mining queries, and index and classify certain types of data to allow more  
13      efficient future query execution, and to also suggest useful information to its  
14      client agents. The answers to popular queries can be priced and cached, so that  
15      providing agents receive value whenever an answer is sold to another agent. The  
16      system of SDI might advertise a set of queries to allow vendors to select pre-  
17      computed results.
- 18      • Matchmaking applications. We described in the top-level description of SDI a  
19      technique to implement “persistent queries”, which reside on the central data  
20      warehouse’s query execution module, and are triggered whenever the correct  
21      conditions exist in the data base. One typical use of such a query is to say  
22      “introduce me to other users with property P”, such that the system introduces  
23      user A\_1 with user A\_2 if A\_1 has the property required by A\_2 and also A\_2 has  
24      the property required by A\_1. This is equivalent to “introduction by mutual  
25      consent”, and is possible within SDI without *information leakage*, because the  
26      only agents that are informed of a match are the agents with the correct properties.  
27      Applications exist to finding a business partner, funding a new start-up (incubator  
28      Co.), forming an interest group, n-way negotiation, introducer system (by mutual  
29      consent).

This section describes specific applications of SDI-based static data mining. A central example is collaborative filtering and personalization applications in electronic commerce, where consumers and vendors provide information to SDI, and the information can be queried within the price and data-access rules placed with the data by owners of the information. We refer to the variation of SDI with consumer profile information as the “iamworthit” system, because users can place profile information in a shared database and receive payments in return for queries performed by vendors. Vendors can use the profile information to build better customization models, and provide customized products to customers based on their profiles and what has worked with other customers with a similar profile.

We limit our attention in this section to “static” datamining, which as defined in the mid-level SDI description allows agents to query the data, but without contacting the agents to request more information. All queries are performed on the basis of the information already submitted to the database, and the rules associated with that information. Within the set of *static* queries, we do allow a vendor to identify a specific agent within a query command, so that if a user provides an identifier to a vendor then the vendor can query the database with its collaborative filtering model and determine an appropriate action to take.

It is possible to request an *action* as the result of a query, for example “send message X to all users with profile information Y”, so long as this message is not contingent on as yet unknown information about the user. We describe specific examples, for example to a smartbrowsing system, an education portal, and an advertising network.

One of the core purposes of SDI is to provide a common location and format for information that has been gathered from a wide variety of sources and that might require different sorts of analysis. Since its framework is designed to handle different types of data and algorithms, SDI can be used as a platform to explore and exploit the rich

connections that potentially exist within and across the databases of different vendors and customers. The system is designed to allow vendors to execute queries over profile information provided by multiple agents, and ensure that all queries are consistent with the policies outline by agents as information is first submitted.

The central SDI server can also support cross-vendor and single-vendor personalization tools, such as multi-attribute collaborative filtering techniques. The queries can be executed to enhance a vendor's model, without providing the raw data to a vendor. A vendor's model can be enhanced *without* explicitly revealing any information about user profiles. The SDI server ensures the integrity of data, and prevents data being used for unauthorized purposes. A query will receive access to data as permitted by certificates presented by the querying agent, and to the extent that the querying agent is willing to make payments defined in the *price-access* rules associated with the data. As discussed in the top-level SDI description, the query-execution module implements an internal market with the data, and executes an agent (i.e. a vendor's) query as cheaply and efficiently as possible.

### 1.1 Statistical Techniques for Multi-dataset Collaborative Filtering

This section describes a statistical method for cross vendor and cross data-set collaborative filtering. The example considers particular types of data and analytical methods and suggests forms of validation that can be made available within SDI.

In describing this system, we show how the Secure Data Interchange architecture can integrate the architecture issued U.S. Patent No. 5,754,939 "System for Customized Identification of Desirable Objects" into a system for secure data exchange between multiple parties. The aforementioned patent teaches a method for profiling objects and users over a bi-directional distributed network, such as: an ISP, multiple ISP networks, a Web hosting network, or server software (such as data mining or recommender software) that is linked to a coalition of sites (such as a portal or Internet mall).

The current invention, the system of Secure Data Interchange, allows correlations to be identified between vendor's data sets, that allows accurate profiling through the application of statistical methods, without providing vendors with explicit access to the profiles of users

12 In the system for SDI we push control of the profile for each user to the client software that  
13 runs on the machine local to the user, and provide for personalization through dynamic  
14 processing of information on the user's client machine. Similarly, we enable vendors to  
15 exchange data sets only to the degree that is mandated by users, and provide technical  
16 solutions to enable significant leverage of data while maintaining user privacy.

17 The supporting architecture as stated in the above referenced patent also allows for profiling  
18 statistics to be collected and processed in a distributed manner. In the present invention the  
19 profile generation capabilities can be implemented at various levels, depending on where  
20 profile information is physically located. As described earlier, the central SDI data  
21 warehouse can be nothing more than a “virtual” database, with multiple links to data that  
22 physically resides on client machines that belong to users and vendors.

23 To enable useful cross-vendor profiling, vendors can submit web pages that are tagged with  
24 profiles of target objects, user quality ratings based upon overall quality as well as other  
25 criteria (e.g., value, price, entertaining, informative graphic/visual appeal, etc.), location data  
26 (for target objects representing physical or geographical items), etc. User information, in  
27 addition to profiles, can include data mining and trend analysis statistics, and user provided  
28 ratings for target objects.

As previously described in this patent, various conditions can be placed on the way in which a set of data may be used (i.e., can the user make a personal copy of the dataset?), as well as on the privacy controls put in place. It might well be that a vendor or a user is willing to



many standard methods that can be used to achieve this, such as Principal Components Analysis. Another approach is to adjust the granularity of the data, if at all possible. In a music store analysis, for example, there might be many more album titles that artists (since each artist can produce multiple albums). In such a case, purchases could be recorded by artist rather than by album, greatly reducing the dimension of the customer vectors' purchase space.

### 1.1.1 Data Structure

In this application there are many types of information which can characterize both users and items. SDI is intended to function as the intermediary between a vast web of vendors, on the one hand, and individual consumers, on the other hand. Major sources of data include:

- 1) Demographic. Such data will most likely be elicited by SDI from vendors and consumers when they initially register for the service, and details very general characteristics about them. It will consist of numbers and categorical values (age, zip code, sex, level of education, etc.).
- 2) Commercial. This is the kind of data any that vendor collects in the course of doing business (especially e-commerce); generally, it links customer codes to purchase items, dates, quantities, and prices. Depending on the nature of the business, this data could be fairly complex, and might well include text. For example, one could imagine that a bookstore, in addition to keeping track of its sales history, collects book reviews, author profiles, and plot summaries.
- 3) Behavioral (vis-à-vis the Internet). A user's client-side SDI proxy can monitor his/her browsing behavior on the World Wide Web, monitoring the pages a user hits, the click stream and content requested, etc. Click stream information can be useful, for example, because it can indicate a user's interest in the information that it is presented with.

We assume that vendors and/or a third-party annotate web pages with *tags*, that provide a commentary of a web page and allow meaning to derived from a user's browsing behavior.



category  $i$  is represented as an  $n_i$  dimensional vector,  $y_i$ . Hence, the total number of dimensions used to describe the full set of  $n$  categories ( $y_1, \dots, y_n$ ) is

$$\sum_{i=1}^n n_i$$

Note that sparse methods are especially useful here, since a categorical vector  $y_i$  will typically consist of mostly zeroes, with a single non-zero coordinate representing the categories' value (i.e., we encode the color red, using the previous example, as (1,0,0) ).

Note also that category vectors with different values are treated as orthogonal by the system.

A final issue is the representation of text. As described in previous related patents, all relevant blocks of text in the database are converted into a dictionary that maps unique strings to the number of times they appear in the database. An appropriate TF/IDF weighting function is chosen and calculated for each of the  $p$  words that appear in the dictionary. The full set of text connected to a single customer can thus be represented as the vector ( $z_1, \dots, z_p$ ), where each  $z_i$  equals the number of times the word  $i$  appears in text related to the particular customer multiplied by the TF/IDF score assigned to word  $i$ .

In summary, when a database describes its customers using a combination of numerical values, categories, and text, customer  $i$  can be represented by the vector  $c_i = (x_1, \dots, x_n, y_1, \dots, y_m, z_1, \dots, z_p)$ .

### 1.1.5 An Example Profile Vector

Suppose we have a database containing information on customers' ages, their musical preferences (i.e. an answer to a survey asking: "Which do you prefer, Mozart or the Beatles?"), and the contents of the emails they've written. Furthermore, suppose the only salient variables in all the emails written consist of the words "Beatles", "Mozart", and "practice", and that we are using the function

$$TF / IDF(x) = \frac{1}{\sqrt{n_x}}$$

Where  $n_x$  represents the number of times word  $x$  appears in the dictionary. We now want to represent one of the customers in the database; he's a 10-year-old boy who prefers Mozart to the Beatles, and who wrote an email to his friend that mostly describes his attempts at



1 practicing Mozart, but in passing mentions his sister's new Beatles CD. Suppose he uses the  
2 word Mozart 2 times (although it appears 456 times in the full database of all customers'  
3 emails), the word Beatles 1 time (appears 217 times in database), and the word practice 3  
4 times (appears 77 times in database).

5 We define the following coordinates:

6  $x_1 = \text{age} = 10$

7  $y_1 = \{\text{Mozart, Beatles}\} = (1, 0)$

8  $z_1 = \# \text{ of times customer uses word "Beatles"} \times \text{TF/IDF}(\text{"Beatles"}) = 1 * 0.067 = 0.067$

9  $z_2 = \# \text{ of times customer uses word "Mozart"} \times \text{TF/IDF}(\text{"Mozart"}) = 2 * 0.047 = 0.094$

10  $z_3 = \# \text{ of times customer uses word "practice"} \times \text{TF/IDF}(\text{"practice"}) = 3 * 0.114 = 0.342$

11

12 *In our example, then, we might encode this boy as customer 1:*

13  $c_1 = (x_1, y_1, z_1, z_2, z_3) = (10, 1, 0, 0.067, 0.094, 0.342)$

#### 14 **1.1.6 Choosing an Appropriate Level of Data Granularity**

15

16 We define the term granularity to denote the level of detail available within a given set of  
17 data, which is often structured hierarchically. Suppose a grocery store database contains  
18 records for a box of flavored gelatin powder. This could be categorized in a variety of ways;  
19 moving from the most specific to the most general, we might treat this data point as "12.5  
20 ounce, strawberry flavor, Jello-brand gelatin dessert" (which would be entirely different  
21 from "12.5 ounce, banana flavor, Jello-brand gelatin dessert"), or as "12.5 ounce Jello  
22 gelatin" (a categorization which would treat as identical the strawberry and banana Jellos),  
23 or as "flavored gelatin", or as "dessert", or as "food", or as "grocery".

24 When analysis is performed on such data, the level of granularity chosen will have a strong  
25 effect on the outcome of the analysis. If the level of granularity is too fine-grained, the data  
26 will be too sparse, although it could be potentially aggregated to the next highest level of  
27 granularity. If the granularity is too coarse, the results of the analysis might be overly  
28 general (e.g., a customer would find a collaborative filter useless if the only  
29 recommendation it makes for a dessert choice is "go to the grocery section of the store").

1 Since the level of granularity will have a salient effect on the outcome of an analysis, it  
2 should be chosen very carefully, and might well play a factor in pricing when a vendor  
3 chooses to sell its data.

#### 4 **1.1.7 Statistical Methods for Data Analysis**

5  
6 In order to perform a wide range of analytical tasks, SDI needs to make use of a variety of  
7 computational approaches. These are described below, starting with the simplest methods  
8 first.

- 9 • (1). Standard Database Searches

10  
11 Since most of the data will be stored in centralized databases, simple searches,  
12 queries, and data filters can be implemented by means of standard SQL commands.  
13 Typically, data will be collected or sorted using efficient database calls before being  
14 fed through analysis routines; once complete, the results can be fed back out to the  
15 database environment for further efficient manipulation.

- 16 • (2) Metrics – Measuring the Similarity Between Profile Vectors

17  
18 Given two customer (or vendor) profiles,  $c_i$  and  $c_j$ , it is frequently desirable to know  
19 how similar they are. For this purpose, we define the similarity metric  $M(c_i, c_j)$  to be a  
20 function that takes as input two customer vectors and returns as output a numerical  
21 value in the range  $[0,1]$ . When two customers  $c_i$  and  $c_j$  are identical,  $M(c_i, c_j)=1$ ; when  
22 they're completely different,  $M(c_i, c_j)=0$ .

23 The problem is somewhat simplified by the fact that we treat all customers as vectors.

$$M(A, B) = \cos \theta = \frac{A \cdot B}{\|A\| \cdot \|B\|}$$

24 Given two customer vectors, we can use the correlation between them to serve as our  
25 metric:

26 Note that  $\theta$  here represents the angle between the vectors  $A$  and  $B$ , and that we expect  
27 all coordinates of the vectors to be positive (in order for  $M(A,B)$  to keep its output in  
28 the range  $[0,1]$ ).

In more complicated cases, however, a customer vector might contain multiple fields with varying ranges of values. For example, we might have customer vectors of the form  $c_i = (\text{age}_i, \text{income}_i)$ , in which the maximum age is 80, but the maximum income is 300,000. In such cases, the coordinates with larger values will dominate the similarity metric, overwhelming any influence that smaller fields might have.

This requires a normalization of the customer vectors, which can be done in several different ways. One approach would be to scale every coordinate by the maximum observed value, forcing all coordinates to lie between 0 and 1 (again, enforcing the rule that all coordinates must be positive).

$$c_i = \left( \frac{\text{age}_i}{\max(\text{age})}, \frac{\text{income}_i}{\max(\text{income})} \right)$$

The only problem with this is that if a coordinate's maximum value is an outlier (being vastly bigger than the typical value), most of the coordinates' values will seem unusually small once they are scaled by the maximum. In such cases, it might be better to scale the values with a "squashing" function such as the sigmoid, which deadens the impact of extreme values; one such configuration would be the following:

$$\begin{aligned} \overline{\text{age}_i} &= \frac{\text{age}_i - \text{mean}(\text{age})}{\sigma_{\text{age}}} \\ \overline{\text{income}_i} &= \frac{\text{income}_i - \text{mean}(\text{income})}{\sigma_{\text{income}}} \\ c_i &= \left( \frac{e^{\overline{\text{age}_i}}}{1 + e^{\overline{\text{age}_i}}}, \frac{e^{\overline{\text{income}_i}}}{1 + e^{\overline{\text{income}_i}}} \right) \end{aligned}$$

Note that the mean and variance of the data points are used to fully normalize them, such that the sigmoid function will spread the values somewhat more evenly between zero and one.

The previous approaches are especially useful for single numerical fields, which might well overwhelm each other if some sort of normalization isn't performed.

1 A different problem arises for text or large categorical fields, since they can  
2 potentially consist of hundreds of coordinates capable of overwhelming the influence  
3 of single numerical fields. Suppose we believe the age of a customer is as important  
4 as the text of articles read. In such a situation, the thousands of coordinates devoted to  
5 the text field would dominate the metric's behavior, negating any influence that age  
6 would have on our measure of similarity – clearly not a good situation.

7 A solution to this would be to find the correlations among the fields taken separately,  
8 then average the result. That is, if each customer  $c_i = (age_i, text_i)$ , where  $text_i$  is a  
9 vector with a very high number of dimensions, we could define the metric:

$$M(c_i, c_j) = \left( \frac{corr(age_i, age_j) + corr(text_i, text_j)}{2} \right)$$

10 Where

$$corr(c_i, c_j) = \frac{c_i \cdot c_j}{\|c_i\| \cdot \|c_j\|}$$

11  
12  
13 The result is a metric that gives equal influence to each field.

14 • (3) Forming Vectors Into Groups

15  
16 The process of classification is essential to collaborative filtering, as it allows  
17 different vectors to be formed into groups based on some measure of similarity. If we  
18 are able to create groups of customer vectors, for example, we can then give  
19 individual customers recommendations based on the patterns of their group-mates,  
20 who presumably have similar tastes.

21 K-means Clustering and Nearest Neighbor algorithms are extremely useful for  
22 grouping purposes: previous iReactor patents give a full and detailed description of  
23 our customized versions. This section gives a brief overview of these methods.

24 (3.1) Clustering

25 K-means Clustering is an algorithm used to partition a coordinate space such that all  
26 vectors in a given partition are more similar to that partition's vector average (the

centroid), than to the centroids of any other partition. It is a process that iterates over the following steps:

0. "Seed" the coordinate space with the initial centroids, which are vectors used to describe the centers of the clusters, in the sense that they are the average of all the vectors currently assigned to the partition. This can be done randomly (assigning centroids random coordinates) if no other information is available, or it can be guided by pre-existing information. For example, if we wish to cluster vectors of music customers, we can use information about musical genres to create initial partitions that correspond to pop, gospel, classical, etc. This will locate the centroids in well-spaced intervals across the coordinate space.

1. Assign vectors to the most similar centroids. This is done for each vector by scanning across all centroids and calculating similarity  $M(\text{vector}, \text{centroid}_i)$ ; once finished, the vector is assigned to the cluster whose centroid has the greatest similarity. In this stage, vectors may switch their allegiance from one centroid to another, if the relative distances to the vector have changed sufficiently since the previous iteration. If no vectors change their allegiance, the iteration process is complete, and the algorithm stops.

2. If the iteration is not complete, recalculate the centroids by setting them equal to the average of those vectors that have been assigned to them. Go back to step 1.

Once the algorithm converges, the vectors are grouped into clusters. The centroids' coordinates as well as the identity of cluster members is useful information that can be passed on to subsequent stages of analysis.

### (3.2) Nearest Neighbor

The nearest neighbor algorithm, simply stated, creates a list of those vectors in a database that most resemble a particular target vector. This is accomplished by comparing the target vector, in turn, to every other vector in the database; the similarity between them is recorded, and once the comparison loop is complete the list of similarities is sorted. The top k members of this list are returned as representing those k vectors which most resemble the target.

- (4) Generalizing Across Databases



1 Suppose we have a set of databases, A, B, ..., Z. Taking each database in turn, we  
2 cluster it using all available data. Thus, using every record in database A, we group  
3 A's customers into clusters  
4  $A_1, A_2, \dots, A_n$ . Taking database B, we create clusters using all of B's information,  
5 creating customer clusters  $B_1, B_2, \dots, B_m$ , and so forth.  
6 Now, scan both databases for common pseudonyms (representing those customers  
7 who have interacted with both vendors under the same pseudonym) and create count  
8 variables  $w_{ij}$  to represent the number of pseudonyms that appear jointly in  $A_i$  and  $B_j$ .  
9 We can now produce the probability that a pseudonym appearing in  $A_i$  will appear in  
10  $B_j$ :

$$P(B_j | A_i) = \frac{P(B_j \wedge A_i)}{P(A_i)} = \frac{w_{ij} / total}{\sum_{j=1}^m w_{ij} / total}$$

$$total = \sum_{i=1}^n \sum_{j=1}^m w_{ij}$$

11 For example, if we have a database of airline ticket purchases and a database of  
12 restaurant visits, we can create clusters, in the first case, of customers who travel to  
13 similar destinations, and in the second case, of customers who eat at similar  
14 restaurants. Given that a particular customer belongs to a cluster of people who  
15 frequent Caribbean restaurants, we can infer which travel packages would most  
16 appeal to him based on the linking probabilities, as defined above.

17 • Multivariate Extensions:

18  
19 If we have a third database C, and there are a large number of pseudonyms common  
20 to A, B, C, the above probabilities can easily be extended. For example, knowing that  
21 a customer appears in  $A_i$  and  $B_j$ , we can calculate the linking probabilities to any  $C_k$ :

22





counted, for TF/IDF purposes, by accompanying dictionaries), the union of the words will define the text coordinates of the new common-information profile. When word counts are being mapped from their original databases to the new vector, the original TF/IDF weightings may be used, or new TF/IDF weightings may be created (using a dictionary constructed from all the databases' text taken together).

Once analysis has been performed, certain common-information profiles will be grouped together by their shared similarities, although the pseudonyms they represent may have been originally drawn from different databases. Such groups will represent links between different databases, and may be used for predictive purposes (see end of example).

### **1.1.8 Example of Cross-database Analysis**

In this example, suppose that the central SDI data warehouse contains data submitted by the following vendors:

- A. A travel agency keeps track of tickets sold, and vacation web pages browsed.
- B. A bookstore keeps track of books sold, and stores an electronic version of the New York Times Review of Books.
- C. A sporting-goods and clothing shop, keeps track of purchase items sold (which includes magazines, for which electronic text exists).

A certain airline wants to promote various vacation packages it has available, which include both European and Caribbean vacations, as well as singles and family packages. Although it has leased rights to databases A,B, and C, it turns out that no customer pseudonyms appear in more than one database at a time – in other words, there are no shared records.

A vacation expert is hired to create a common-information profile. He creates the following information vector: (list of tropical countries, list of European countries, family score, list of sports, text)

Note that the family score is a numerical value ranging from 1 (young singles) to 10 (many small children), and indicates what kind of person the customer is (a party-oriented student vs. a sedate father of three).

The expert the creates the following mappings:

- 1       A. Travel Agency. Link destinations of tickets sold to country fields (i.e., the number  
2       of trips to Germany by a customer would be placed in the Germany field of the  
3       common-information profile). Link sales of children's tickets, or requests for  
4       children's meals, to family score. Put web-page data into text field.
- 5       B. Bookstore. Link travel books' text to country lists. For all books purchased by a  
6       customer, map text from book reviews into text field.
- 7       C. Sporting-Goods store. Map warm-weather clothing (and swim gear) to tropical  
8       countries, ski gear to countries with skiing areas. Map sales of toys or children's  
9       clothing to high-value family scores, map revealing-bikini and student-discount  
10      sales to low-value family scores. Map text from magazines purchased by a  
11      customer to text field.

12  
13      These mappings are then applied to each database, generating a full set of common  
14      information profiles. These are then clustered, forming groups that share commonalities.

15      The expert can now do several things with the results. First of all, he identifies the general  
16      "flavor" of each cluster (e.g., families with small children that enjoy winter, Europe, and  
17      skiing); the pseudonyms contained within each cluster can then be targeted for vacation  
18      packages suitable to their tastes. Secondly, the fact that pseudonyms from different  
19      databases have been clustered together allows the expert to plan cross-category marketing. If  
20      certain travel-book-buying parents have been grouped together with parents who bought  
21      their children swimsuits and scuba toys, it may be that they share a preference for family  
22      activities that take place in warm places or by the seashore. Hence, the book-users might be  
23      advertised various ocean-related sports goods appropriate for young families, and likewise  
24      the swimsuit-users might enjoy getting recommendations for travel books that describe  
25      tropical destinations that are especially fun for children. That is, if the goal is to cross-  
26      market items from A to customers in C, the most logical source of recommendations would  
27      be the people in A who have been grouped with the people in C.

### 28      **1.1.9 Methods for Validation**

29  
30      To a large degree, the overall success of an SDI analysis is the relevance of the connections  
31      that are inferred from the data. It is often the case that a certain amount of validation is



adding it to a running total. The system with the highest total can thus be judged the most effective, since it most strongly recommended items that the customers did, in fact, end up purchasing.

Because the result of this type of validation is a quantitative score, it is possible to automate the model selection process. Given a set of analytical approaches (each with its own array of parameter settings), it is possible to loop through the full parameter space (using a grid of evenly spaced numerical values, if needed, to reduce dimensionality), computing a validation score at each iteration. Those combinations of algorithms and parameter settings that demonstrate the best performance could be chosen as the top candidates for the final system configuration, since they do the best job at predicting customer behaviors.

#### (1b) Quantitative Approaches – Dynamic Method

The problem with the hold-out approach to validation is that it isn't dynamic, since it doesn't reflect the impact that the recommendation system has on the customers once it is implemented, and may be based on data that doesn't contain current trends. After all, it is better to predict what the customer will buy rather than what the customer has bought in the past.

A better approach is to run a controlled experiment against the actual customer base. First, the pool of customers is split at random into different segments. Next, each approach under consideration is used exclusively to make predictions for a given segment. Once the trial period is over, each system is given a score based on how valuable its recommendations turned out to be (this could be measured by total sales generated, for example, or by the number of times a customer made use of a recommendation).

#### (2) Human Expert in the Loop

Although quantitative methods can automate the validation process to some degree, at the beginning of many projects there is so much raw input data available and so many decisions that have to be made about the analytical approach that an automated process would have to test a prohibitive number of combinations of data, algorithms, and parameter settings to get optimal results. In such cases, it is useful to employ a



1 Iamworthit allows individuals to receive payments for the information that is collected by an  
2 SDI client proxy and stored in the central SDI data warehouse, for data mining purposes.  
3 Iamworthit also allows individuals to receive payments in one-to-one interactions with  
4 vendors in return for providing vendors with information that allows them to make an  
5 appropriate offer. Secure data interchange supports the useful exchange of information  
6 between agents without allowing vendors to collect and distribute information about users  
7 without the knowledge or permission of users, as is possible now via the system of cookies  
8 and affiliate networks, such as that operated by [www.doubleclick.com](http://www.doubleclick.com). In the system of  
9 Iamworthit vendors benefit through well-targeted advertising (both push and pull), the  
10 ability to customize information and services (even to first-time customers), and access to a  
11 large database of information about buyer purchasing habits. Individuals still release profile  
12 information to the central SDI database for the purposes of controlled access by vendors,  
13 and also provide vendors with profile information during interactions and allow vendors to  
14 execute queries on information and receive the benefits of personalization without directly  
15 accessing the information. Users can maintain multiple pseudonyms and profiles, but within  
16 SDI vendors can still access information across pseudonyms and use cross-web (broad) and  
17 single-vendor (deep) information to build robust models of buyer behavior. Buyers benefit  
18 through personalization with privacy, and financial rewards in return for releasing profile  
19 information.

20 Within B2C e-commerce, secure data interchange can also support a system of *time-of-*  
21 *purchase competition*, which allows an individual to use the profile management capability  
22 of SDI to provide vendors with information about a user at the time of purchase, and allow  
23 competing vendors to offer the same product at a better price, or a better product (for the  
24 user) at a good price. With time of purchase small entrants to the marketplace can make  
25 counteroffers to users that are about to purchase a product or service from another vendor,  
26 and can compete in small parts of the marketplace without investing heavily in advertising  
27 and brand awareness. Time-of-purchase requests can be made by user clients to iamworthit,  
28 to request that iamworthit cascades purchase requests onto other vendors, collects responses,  
29 and then return them to user clients. This extension of SDI is discussed in Section 2.1.

30 As a commercial strategy, one might make it a necessary condition of belonging to  
31 Iamworthit that client machines submit profile information to the SDI central data

warehouse, so that vendors can perform useful data mining and then provide customized products and information to users, for example based on models of collaborative filtering.

### 1.3.1 System Overview

Client machines in Iamworthit implement an SDI proxy on-top of a user's regular web browser. The role of the proxy is:

- a) Profile and identity management as the individual interacts with server computers of on-line vendors.
- b) Collect and manage profile information, with information periodically submitted to the central SDI data warehouse.
- c) Release profile information to vendors according to profile management policies.

In Iamworthit the client-side SDI proxy is configured by users to periodically push information collected about the user to the central SDI data server, i.e. click stream data, profile information, purchases made, information requested, etc. Profile information is associated with price-rules, as in the top-level description of SDI, and can also be adjusted with random perturbations to protect the identity of a user under multiple pseudonyms.

The client periodically sends update messages to the database, for example with information about new activity (e.g. web browsing, purchases, adverts not selected, etc.) SDI allows the information to be used by vendors to build models that enable good personalization of products, services and adverts.

Users might choose to randomize profile information to protect sensitive information and to prevent vendors linking users across different pseudonyms. Clients submit perturbed data points and the distribution that was used to generate the noise (see the section on random noise perturbation).

### 1.3.2 Privacy-protected data mining

Vendors can gain access to profile information in the central SDI data warehouse, as long as they meet the criteria of the agents that submit information on behalf of users, and pay the price of data access. The profile information can allow vendors to build complex models for personalization, negotiation, and advertising, based on information about purchases made by different types of users.

1 Users submit profile information to a central database that can be used for profiling, without  
2 revealing their identities. Client-side SDI proxies track user activity on-line, across multiple  
3 sites, and submits data, randomized if necessary, and only according to a user's preferences  
4 to the central SDI data server. The data is useful for building models of buyers, for example  
5 purchasing patterns, for the purposes of personalization of information and adverts. Agents  
6 that submit information retain ownership of the data.

7 Simple data mining queries include:

8 (a) Compute the average income level of people purchasing camcorders. Suppose  
9 that a vendor has pseudonymous identities of its recent customers, and wants to  
10 compute their average salary. It does not know the salary of any of the customers,  
11 and individual customers will not release their salary to the vendor. However, if  
12 some of those customers have provided information about their salary to the  
13 central SDI data warehouse then the vendor can compute the result to its query.  
14 Agents might associate price rules with information about their salary that allow a  
15 vendor to query that information so long as the identity of the agent is not  
16 revealed, i.e. so long as the salary is provided anonymously. We described earlier  
17 in the patent description how price rules can define different prices for different  
18 types of information access.

19  
20 (b) Compute the total donation to a fund for bone cancer research. Individuals do not  
21 wish to release information about their specific charitable donations, but might be  
22 happy to release that information anonymously. Again, a query to compute an  
23 average donation can be formulated and executed in the central SDI data  
24 warehouse.

25  
26  
27 (c) Perform collaborative filtering across multiple fields. This is explained in more  
28 detail below, and is possible within SDI because a general query can be  
29 performed so long as the querying agent makes payments for data access in  
30 accordance with the price rules of agents which submit information to the shared  
31 database. Note also that randomly perturbed data does not prevent the



computation of correlations between fields, so long as the random perturbation is “small” with respect to the value of the data record.

### 1.3.3 Importing Off-line Data into the Data warehouse

The central SDI server can associate off-line information about a user with a user's on line pseudonymous profile, even though the central server does not know the user's pseudonym IDs. This can only be done with the user's consent, and may also involve appropriate compensation. Within the system of iamworthit we can credit users for both off-line and on-line information.

Merging a marketing database with SDI user profiles can be useful both to initialize the database, for example when asking a user questions to generate an accurate user profile rapidly and efficiently. Off-line data can also add useful richness to on-line profiling information, which may be largely contextual and low on details/factual information. For example, off-line data can include information such as whether a user owns a car, rents an apartment, has house insurance, life insurance etc. SDI can also extrapolate correlations to other user profiles, on the basis of common SDI-profiles, for example using statistical techniques.

It is often the case that individual customers appear in some databases, but not in others. Under normal circumstances, an analyst working across different databases would be faced with a large number of incomplete customer records, each with gaps corresponding to the fields of the databases to which they don't belong. A solution to this problem is offered by SDI, which is capable of drawing correlations between different databases - this information can be used to generate predictions to fill in the gaps of incomplete customer records. The result is a full set of customer records that can be meaningfully sorted or filtered by any of the combined fields, and which can now be handled as a unified set of data, suitable for use by standard database analysis systems.

In a typical example, SDI might be used to combine a demographic database, such as the one offered by the Econometrics Corporation, with a commercial database, such as the one offered by Claritas. The Econometrics database consists of 180 million different customer records, but at a fairly coarse-grained level of detail, consisting of such information as age, gender, family status, location (at the state, city, or zip code level), and personal income. In

1 comparison, Claritas offers a smaller base of customers, but includes information of  
2 arguably higher quality, since it breaks customers down to the geocode (sub-neighborhood)  
3 level, and includes much more detailed information on personal spending habits across  
4 hundreds of different purchase categories. A logical reason to combine these databases  
5 would be to supplement information about customers in the vastly broader demographics  
6 dataset with particular predictions about their personal preferences and likely commercial  
7 spending habits. One could imagine using this augmented data set to support a web site that  
8 instantly customizes itself to new visitors' preferences. Since the number of records in the  
9 Econometrics database is equivalent to roughly 72% of the population of the United States,  
10 it is likely that most first-time visitors to the site will already have a "thumbnail sketch" in  
11 the system, and can thus be greeted with an page appropriately configured to their personal  
12 tastes.

13 The technical details of the combination process (which have been described elsewhere in  
14 the patent) to a large degree depend on the amount of overlap between the databases, that is,  
15 the number of customer records which are shared in common.

16 Suppose the demographic databases' fields are coded  $(x_1, \dots, x_n)$ , and the commercial  
17 databases' fields are coded  $(y_1, \dots, y_n)$ . Suppose further that customers in set A appear only in  
18 the demographic database, customers in set B appear only in the commercial database, and  
19 customers in set C appear in both.

20 The process of supplementing the fields of customers A depends completely on the  
21 derivation of the distribution  $f(y_1, \dots, y_n \mid x_1, \dots, x_n)$ , which describes the correlation of fields  
22 in the commercial database on fields in the demographic database. As previously discussed  
23 in the patent, different techniques may be used to create this distribution, depending on the  
24 size and variety of C.

25 As a concrete example, one could imagine that set C includes customers from rural areas.  
26 The demographic database would reveal that, although their incomes aren't huge relative to  
27 the national average, they tend to spend a lot of it (i.e. are active consumers), have large  
28 families, and purchase large vehicles. The commercial database might show that they  
29 enjoying hunting magazines and Ford trucks. If they live inland, they buy hunting  
30 equipment, if they live near the ocean, fishing equipment.

1 If these trends are dominant in set C, they will impact the distribution function. Thus, when  
2 a browser from a small town in Texas with a typical income pattern visits the automated  
3 website, he could be greeted with discounts on truck accessories and a small sidebar with  
4 news on the hunting season. On the other hand, a visitor from a small town in Maine might  
5 be given the same truck discounts, but would have news on the fishing season.

6 Although the demographic dataset is arguably the weaker of the two in terms of content, the  
7 fact that it contains even a small amount of information on most people in America makes it  
8 very valuable for handling first-time visitors, since most of them will appear in it. By using  
9 SDI to leverage the more detailed information in the commercial database, we are able to  
10 supplement the rough demographic data with predicted commercial preferences. This allows  
11 us to construct more detailed thumbnail sketches for each customer, allowing our reception  
12 of first-time visitors to be much more appropriate (since knowing personal hobbies or  
13 interests tells us much more about a person than general income level).

### 14 1.3 Static Query Execution: Central Applications

15  
16 Some key applications of SDI are:

- 17  
18 • Assessing the Value of Data. Plug together sets of data, and measure predictive  
19 accuracy.
- 20  
21 • Matching Data Across Vendors. Find patterns in common pseudonyms, denoting  
22 common areas of interest; use catalogues of order codes and item description to  
23 find similarities across data sets.
- 24  
25 • Targeted Recommendations; e.g. match customers to their nearest neighbors in a  
26 data set and generate recommendations for users, collaborative-filtering style  
27 application.
- 28  
29 • Leveraging Portal Data. Use data from portal to leverage data needs for ISP  
30

- 1. Analyzing Affinities. Suppose a vendor has a list of customers, and knows to  
2 some degree what web pages they visited after leaving vendor site. A large  
3 collection of customers taken from an ISP will contain their web-surfing behavior.  
4 Cluster web sites and cluster customers, finding cluster-to-cluster interactions.  
5 Use this information to classify vendor's customers; gives vendor an edge in  
6 knowing customers' tastes.

8 There are a number of search-based applications, where SDI searches for appropriate  
9 profiles and then requests that SDI makes contact with the users pseudonymously, i.e.  
10 without the vendor receiving any useful information about a user's identity. The contact,  
11 interaction, and business relationship with the vendor occurs under terms of complete buyer  
12 pseudonymity. In accordance with the parent patent application [INSERT US PATENT  
13 No.] the pseudonymous communication may be either email, real-time text  
14 communications, voice (such as the pseudonymous telephony or Internet telephony). In the  
15 case of pseudonymous telephony, instead of a one-time or persistent pseudonymous buyer  
16 address, pseudonymous buyer telephone numbers may be used for the third party to reach  
17 the buyer under his/her terms. Example applications include:

(i) Financial Advice and Financial Planning Services. Often buyers are quite sensitive about the confidentiality of the release of this type of information related to personal financial matters and particularly with certain matters (and perhaps in general) prefer that their financial advisors were unaware of their true identities. Similarly, investment advice or sales communications by stock brokers are another application where similar buyer information is typically disclosed.

(ii) Insurance Agents & Brokers. For many types of insurance, (e.g. health, life, casualty) personally sensitive information is disclosed by buyers to their agents and brokers. Initially, before insurance services are purchased, it is possible that useful detailed quotes and/or insurance advice could be provided to a buyer pseudonymously.

(iii) Legal advisors. There are a variety of legal disciplines in which the associated legal services delve into highly sensitive personal information (e. g., bankruptcy law, divorce law, criminal law, etc.) Many lawyers also offer to first-time prospective clients a free consult in which such a privacy-enhanced communications system could be initially beneficial to the parties.

(iv) Family Counseling and Psychological Counseling. The parent patent application also suggests these applications which often involve the exchange of highly confidential personal information.

(v) Medical Consultations Involving Drug Prescriptions. In this variation, although medical consultation may be conducted anonymously, in order for a physician to prescribe medication, the identity of the patient must be known (within the current regulatory legal requirements).

(vi) Advertising network (where the clustering or nearest neighbor algorithm interact with the ad server). Ads on that ad server's database, e.g., on the sites most frequently visited by the user are periodically uploaded to SDI such that the ad(s) of highest predicted interest to the user is presented upon the user visiting that site. Preferably the target object profiles of all the ads on the ad server database, as well as the profiling algorithm, which is used by the ad server is properly integrated with that of SDI so that it is possible to convert the usage statistics of the data model used by the ad network (which may include ad server data for users collected across the ad networks) into useful statistics for SDI.

(vii) Content sites, e.g., for personalizing news articles (which are again uploaded to the SDI server each news day if those articles would not otherwise be available on that web site. If they are, as with any web content format which is universally accessible, the content may be frequently profiled following retrieval, e.g., by a simple web crawler script.

[illegible]

2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31

16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31

30  
31

1 or in combination as a distributed process upon the caching server in  
2 accordance with the method taught by "Broadcast System for Reduced  
3 Memory Devices and Asymmetric Networks". (check exact title of this  
4 patent) As taught in this patent (in accordance with the current applications)  
5 the pages predicted to be accessed by the user on the following day are  
6 precached in advance.

7  
8 Typically sites which tend to be visited most frequently are assessed for  
9 selections which are personally relevant but have never been previously  
10 accessed. These would be presented as recommendations to the user if/when  
11 he/she visits that site (which is also probabilistically determined). The fact  
12 that the recommendation (whether generated by the site itself or modified at  
13 the client level proxy), the user's behavior is positively being reinforced by  
14 the fact that the ranking of personal recommendations is prioritized in direct  
15 relation with the prioritization scheme for precaching, i.e. personalized  
16 recommendations and precached pages are substantially identical and as a  
17 result of precaching to the server, or even the client, can be accessed by the  
18 user with little or no latency (thus in the ideal embodiment it may even be  
19 advantageous to highlight the links which have been precached).

20 In a preferred commercial model, the techniques which are above described,  
21 that meta data containing profile information on the pages are encrypted such  
22 that a cache engine that is used to precache cannot be decrypted and read by  
23 certain caching engines which do not possess a desired business relationship  
24 with iamworthit. In another implementation, other competitive intermediary  
25 services to iamworthit may not be able to decrypt and read these web page  
26 profiles in order to present personalized information to the user (even if the  
27 competing protocol were pursued, the accuracy of the iamworthit version  
28 would be substantially more accurate due to its ability to leverage vendor  
29 centric SDI data from the host level proxy.

30  
31 (xi) Auto Insurance Application





numeric values associated with various medical tests, results for which are a numeric value. This data may be of relevance to pharmaceutical companies, alternative medicine vendor and clinics insurance companies hospitals physicians, clinics and home health care providers, the latter three of which may wish to advertise to patient prospects and extend their medical practices. The privacy architecture herein provided is a critical component for enabling access to user data by these commercial entities and is perhaps critical to the extent that users must explicitly authorize the transfer of medical records from either the physical copying or electronic duplication of an associated transfer of such information to an intermediary (SDI) which is trusted by the user.

(xiii) Medical information, such as medical conditions, medical history, active prescriptions, drug reactions, family history, possibly even genetic predispositions (from a genetic profile). Medical insurance information may also be potentially useful for a prospective qualified accessory to be able to readily access in case of an emergency.

(xiv) **Physical location information**—Users or advertisers could, for example: a)

Query a pseudonymous user database to access profiles that are in close physical proximity and match certain criteria, e.g. live in a certain geographical region, had recently attended a meeting or event (or is planning to attend a particular event) had recently communicated with a friend or associate. In another variation, a user could for example, submit a query pertaining to every user in a particular physical space, e. g., a room, hotel or convention center, e. g., identify all users present here who attended Internet World, 1995.

### 1.3.1 Buyer Infomediary

The central SDI data warehouse can be used by vendors that provide services of *buyer infomediaries*, for example providing buyers in B2C e-commerce applications with historical information about previous purchases of users. At present companies such as [www.priceline.com](http://www.priceline.com) make profits because many individuals post take-it-or-leave-it buy offers for goods that are above the reservation price that vendors are prepared to accept for a service. An information infomediary, based around information submitted by users about previous successful or unsuccessful bids would provide for a more efficient marketplace.







1 experimented with such that optimal click-through or transactions occur particularly for key  
2 pages or purchasables. Adding to this statistical data SDI user profile information is able  
3 to reveal (via data mining) much more robust relationships between pages as these  
4 relationships change with user profile features. Accordingly appropriate rules may be  
5 provided for individuals based upon this general user profile information which is gleaned  
6 from interests and behavior before visiting the site, i.e., either from the user-centric or  
7 vendor-centric SDI (while the user visits the vendor's site or alternatively accesses an  
8 affiliate site).

9 In another variation of SDI we can highlight content on web pages with information that  
10 might be relevant to a user, even when the source of the web page is not personalized. This  
11 can be done via collaborative filtering techniques, which might bring in feedback and  
12 comments from other similar users within SDI that are stored in the central SDI database.

13 The iamworthit (user side) SDI database in cooperation with the vendor centric SDI  
14 service can sell to the vendor centric SDI service or other industry or market research  
15 organizations strategic information about the comprehensive behavior activities and user  
16 profiles of visitors and customers of these vendors (as is suggested earlier in the spec).  
17 Additionally, targeted survey questions may be presented on behalf of these entities in  
18 order to extract further information which may be correlated with certain features and  
19 attributes of these users. One such method, rapid profiling is detailed in the parent  
20 application [INSERT US PATENT NO].

21  
22 Vendors can utilize SDI's data mining interface to observe content and product  
23 consumption affinities based on user profiles describing Web wide behavior. This  
24 interface may also enable vendors to observe comparative click through and/or  
25 transaction rates of their competitors, how these criteria are affected by user profile  
26 attributes including geographic criteria (if relevant to that vendor) as well as how these  
27 criteria, (including even individual multivendor customers) are affected by various types  
28 of recommender and rule based engines and further the particular rules and weighted  
29 feature correlations used in generating these recommendations (again as described above,  
30 these rules and feature relationships may be derived automatically through a core sample  
31 of users, who are iamworthit subscribers). A vendor can use statistical data to tune



1 provider can query the data interchange for information about the user. The Internet content  
2 provider sells well-directed advertisements to vendors. Secondly, the data interchange could  
3 sell or rent data to an advertising agency directly, providing information in real-time to  
4 enable the advertising agency to provide more focus in its banner ads for its clients. "Per-  
5 transaction" pricing is a very powerful pricing model that is enabled with on-line banner  
6 ads. It is simple to monitor the number of click-through that are received at a particular  
7 banner, in response to an advertisement. In the off-line world pricing must be based on the  
8 number of impressions, or worst still, the number of mailings sent and it is more critical to  
9 understand the expected value of a campaign up front.

10 The proxy server could also act as an "ad network" itself, and sell focused advertisements  
11 for vendors, and purchase ad-space on the sites of content providers. The on-line domain  
12 provides this unique opportunity for quick experimentation with advertising strategies in  
13 order to get feedback on the likely utility of untested approaches. The system can use a  
14 hierarchical cluster tree to identify the most revealing items in a dynamically responsive  
15 fashion such that the profiles of all of the selections can be generated with the most minimal  
16 amount of interactions with the user (see "Rapid Profiling" section in issued patent entitled  
17 "System & Method for Customized Electronic Identification of Desirable Objects). Thus a  
18 more robust statistical model across multiple vendors is established as a result of the user's  
19 click through response of these intelligently selected virtual banners as well as other pages  
20 which are subsequently navigated through once the remote site is accessed via the banner.

21 In the preferred approach rapid profiling not only dynamically identifies and presents items  
22 which are most revealing of the other items in the collection, it also selects the users whose  
23 profiles suggest the greatest familiarity with these items (i.e., potentially correlated items).  
24 Furthermore, if the system's objective is to find new users or users who may be interested in  
25 the present vendor's other products, products for which little is known, then it will match  
26 users who are least familiar with exemplar items. The idea is to reveal the most significant  
27 data about the user profile with respect to the present collection of items of interest. Finally,  
28 rapid profiling can use direct explicit queries to determine interest on an item(s) or to collect  
29 demographic data on a user.

30 The target object profiles of advertisements on the ad server are matched against the user  
31 profile in order to automatically present the most relevant recommendation(s). Typically,

In another variation, if the data to be disclosed to the vendor is acceptable to the original vendor but she/he is untrusting of the vendor, the data is received by the host-level proxy (another neutral third party) instead of the vendor, thus providing the disclosing vendor with an additional level of security, assurance about the use of his/her data while enabling the users of such a site to access all of the merchandise or content in a completely personalized fashion. Thus these XML tags are stored in association with, but on a separate server from the actual HTML pages stored on the vendor's site. Additionally, these profiles are constantly updated by user profile data conveyed to the host-level server which operates in distributed fashion.

Advertisers may also stipulate additional rules relating to the delivery of ads to buyers in conjunction with performance of certain explicit actions e.g. submission of certain search terms, visitation of pages with certain key words (as a general tendency or on a per-page basis), visitation of certain specific pages or sites, performance of certain location specific behavior **the delivery of which could be performed in accordance with the real-time behavior by the user if desired** (e.g. visitation to a competitor's retail outlet or traveling within the proximity of the vendor's own store), otherwise the presently suggested algorithms may automate most of these manual functions by virtue of its recommendation capability e.g. matching ads to the most appropriate list of buyers or matching ads to metrically similar Web sites which may appear whenever the buyer visits that site(s).

Ads may also take the form of *digital coupons* [KRJM 98] which are priced by vendors according to behavioral information which the buyer is willing to disclose to this vendor.

31 Because many buyers are adverse to viewing additional advertising which is “imposed” in



1 the preferred embodiment, any inserted ads delivered by iamworthit can be “turned off” by  
2 the buyer voluntarily unless the in kind value which the buyer receives in exchange for the  
3 advertising mandates a certain minimum ad impression delivery. Also in light of the above  
4 buyer concerns in the preferred embodiment, there are no additional impressions added to  
5 the buyer’s web browsing experience.

6 Ad blocking technology can be deployed to block existing banners and replace them with  
7 iamworthit ads. Ad blocking is commonly available, and the techniques used are well  
8 known in the art. The HTML source for the ad banners may be either called-up from the  
9 remote server (such as with an ad network) or alternatively from an ad server (typically  
10 purchased by the vendor) on the vendor’s own host server. In either case, the ad blocking  
11 software typically recognize the HTML source which originates from the ad server. A  
12 directory of the HTML source for the various ad servers is maintained with the ad blocking  
13 software (and presumably updated e.g. if new servers are added or change their IP  
14 addresses). The location/physical dimensions of the portion of the page occupied by the  
15 HTML source which the software removes (its “footprint”), may be readily replaced by  
16 another replacement banner (by iamworthit) from wherever the proxy server resides  
17 (typically on the client but potentially on the network). As a result of a potential difficulty  
18 in which the ad server deliberately changes its IP address to avoid recognition, which occurs  
19 constantly and dynamically, it may be possible to recognize portions of the page which  
20 contain image “features” of ad banners, which are inserted from another HTML source  
21 where that HTML source is linked to known vendor sites (which are also linked similarly  
22 from other site’s ad servers and/or are known via their HTML sources to have previously  
23 utilized ad server technology and a “new” HTML source appears instead, etc.

24 This model may be extended to other media domains e.g. replacing digital TV commercials  
25 instead with targeted ads delivered and precached via iamworthit for insertion at appropriate  
26 times i.e. during commercial breaks during standard video programming. Another variation  
27 could be tailored to pre-loading iamworthit advertising to automobiles for insertion in place  
28 of traditional radio commercials. In accordance with the parent patent application, (and as  
29 suggested above), email may be a useful targeted ad delivery medium as well. Per the  
30 patent case, the system is able to classify email (according to its source and content using  
31 implicit or explicit actions of the buyer). Based upon certain desirable confidence threshold





1 following description), user profile which is “most like” the profile of the cluster to which it  
2 belongs, perhaps a median metric.

3 The Platform for Privacy Preferences (P3P) [CR 98; RC 99] provides for the ability to  
4 utilize XML meta-tags to annotate Web pages, and within a system such as SDI we can  
5 allow comments from previous visitors to a page to add annotations. Users can receive  
6 annotations from the SDI data warehouse that are associated with annotations provided by  
7 other users with similar profiles. One of the divisional applications of the parent case  
8 “System for Customized Electronic Identification of Desirable Objects” [Herz 98b] relating  
9 to the automatic creation of virtual communities suggests that users may be automatically  
10 assigned to particular communities (e.g. chat groups, forums, etc.) for this purpose.

11 We might allow users to rate the annotations in the pages, allow SDI to learn how useful  
12 annotations provided by certain users are to new users, and classify users as “experts” that  
13 receive priority in the position of their current and future comments. Future comments from  
14 users with a poor rating history for a particular content cluster may be deleted. A persistent  
15 interface feature on the tool bar or side bar can provide for annotations to also be accessed  
16 by users selecting certain profile features of users as they browse from page to page; for  
17 example a user could identify the comments of a news article about abortion by users who  
18 are self identified as advocates of the Women’s Rights Movement, ultra conservative senior  
19 citizens, teen women or those with a strong interest in alternative medicine or the Catholic  
20 Church.

21 The parent case [Herz 98b] further suggests that users may actively provide ratings in a  
22 completely privacy protected manner according to various criteria of pages they browse.  
23 Profile-based clustering of these ratings and annotations can allow a user to submit as a  
24 query a user profile, to receive a page that rates or annotates a page, and a listing of highly  
25 relevant pages to a particular type of user. This could also be used in a “reverse engineering”  
26 sense, for example requesting the exemplar user profiles for users that visited and gave high  
27 ratings to a particular site, and then using those profiles to find similar sites of interest, or  
28 finding the user features that are in the exemplar profile to understand something about the  
29 content of a web page. The browser interface may automatically display the related links  
30 that are determined dynamically to be most relevant to a user, again statistically estimated

1 via the data in SDI of users' browsing habits, and/or via active page recommendations or as  
2 book marks by those users as being of particular relevance or similarity to the present page.  
3 Browsing methods may include search-based browsing, and also browsing via a hierarchical  
4 navigation menu system, with users classified according to their behavior patterns and/or  
5 ratings which have been actively submitted. The parent case [Herz 98b] also suggests the  
6 use of hierarchical clustering for products, to help in a smart shopping system. Similar  
7 products can be presented together where the feature criteria for creation of the hierarchical  
8 cluster tree could be price or other criteria. In addition to allowing users to view item  
9 selections according to desired selection criteria, a comparison shopping function can also  
10 allow a user to view the attributes of buyers that tend to buy certain items to help to add  
11 confidence to a user that he/she is getting the right product.

12 The above description also describes the use of a hierarchical menu through which groups of  
13 users may be identified by their profile features (wherein a profile feature could even be a  
14 rating criteria itself of for example, an opinion via a site survey). These features could be  
15 used to either selectively filter-out content which falls outside of that criteria as the user  
16 navigates the information or identify if/when pages encountered where these user rating  
17 features are present, thus displaying this user statistical information in conjunction with the  
18 ratings statistics and/or associated annotations if desired.

19 We can allow the user to use one or more organizations to provide features and annotations  
20 to help in a search and information interpretation process, to add appropriate filter and bias  
21 to information presentation as a user browses the Web. Similarly, a user might adopt the  
22 endorsements of a friend for annotations, and editorial content about particular pages.

23  
24 The availability of feature information about individuals that provide high ratings for  
25 particular Web pages can also be useful for vendors that wish to find appropriate  
26 locations for their advertisements. For example, given an "ideal profile" of a consumer  
27 for a particular product, a web page/type of information can be determined that will be  
28 useful to find appropriate customers. This is an alternative method to providing adverts to  
29 particular individuals, instead choosing to provide adverts to limited information domains  
30 that attract appropriate users. Vendors may request additional information, such as site-  
31 specific page view correlations (including time spent viewing each page) in accordance

4

We can also profile users with context information, for example based on their assumed goals (social, business, personal, etc.), information which can be determined by the nature of their current activities.

8 In an interesting extension we can allow users to contact other users that are browsing  
9 similar material with similar profiles in real time, allowing ad-hoc virtual work group  
10 formation. Of course, all of this is done within the carefully managed profile and privacy  
11 managed systems of SDI. The location of a user could play into decisions about physical  
12 meetings. Typically some users may wish to enter into a present (or future) dialogue, which  
13 can be scheduled via calendar agents, which perform automatically scheduled compatibility  
14 meeting/introduction functions.

It is also reasonable to provide the technique of collaborative filtering to identify users whose profiles are particularly similar to the material which is being viewed, e.g., those users who are determined to have a high proficiency level with regards to that particular material. Because the requesting user may be seeking to query the expert or seek edification on that material, typically same consideration is conveyed in exchange for that information.

20 In a variation, a more ----- data exchange may be provided if collaborative filtering  
21 identifies fairly similar users to the presently viewed material which also demonstrate, via  
22 their profiles, considerable complementarity in the particular knowledge which they are  
23 likely to possess. Vendors may also ----- and facilitate this type of  
24 information exchange on their site between these complementary, (yet metrically similar)  
25 customers.

26

In a consumer-vendor commerce application of this form of interactive browsing we can allow users to enter into dialogues with vendors based on their profiles, and share experiences with vendors, regarding their experiences. This can be done in real-time, or statically within an iaworthit style system architecture, where a user receives financial incentives for providing information that is useful to vendors.

1 The parent case [Herz 98b] suggests that *organizations (e.g. vendors, etc)* may be rated by  
2 multiple users across and according to a variety of relevant attributes. One useful extension  
3 is to group descriptive opinions of users according to their profiles, certainly users with a  
4 particular type of profile are likely to have affinities towards certain sites and certain  
5 content. It can be useful to example how the assessment of user's varies according on their  
6 profile information.

7 Another variation is to suggest a *three-dimensional interface* to allow useful diagrammatic  
8 representations of the information and meta-information that a user is presented with. The  
9 parent case [Herz 98b] describes a method to reduce the multidimensional characteristics of  
10 feature and annotation information to only 3 dimensions, and present that information  
11 graphically, via a technique of principle-component analysis. The profiles of the target  
12 objects associated with a user's preferences are represented as points in 3-D space within a  
13 navigational interface. Nodes may also be displayed according to category and sub-category.  
14 Views may also be restricted to user's with the same goals, i.e. business users, social users,  
15 etc. this mind-set can be inferred by the sites that a user is visiting, for example.

### 16 **1.3.9 Smart Interface to On-line Discussions.**

17  
18 The present description provides a methodology by which it is possible to apply techniques  
19 of statistical NLP for purposes of enabling users of live discussion groups and bulletin  
20 boards to dynamically visualize concept (via automatically generated labels short summaries  
21 or even visual representations of the concepts). The parent patent [Herz 98b] describes a  
22 browsing implementation in which it is possible for a user to view an on-line 3-D layout of a  
23 "virtual mall" in which the multi-dimensional space of the object profiles of the items are  
24 collapsed into three dimensions. Further, the display may be customized to the user's  
25 profile by using as the vantage point the position of the user's profile within the space.

26 We can use information retrieval, pattern matching, and collaborative filtering techniques to  
27 refine information based upon the particular relevance feedback from those particular  
28 individual engaged in a discussion. The novel interface enables the extraction of text  
29 segments from all previous discussions. These (multi-discussion) segments are essentially  
30 grouped into clusters and the clusters are representative of "concepts". The statistical  
31 framework is further used to determine what clusters (concepts) tend to follow in sequence

1 what other clusters. Because conceptual flow in human language is quite complex and  
2 perhaps somewhat relative to the individual, it is useful to provide different interface  
3 settings, e.g., for example one which is tailored to the individual (i.e., the layout of the  
4 conceptual flow based upon statistics taken from all individuals which are generally similar  
5 to the interest profile of the present user what it would be for certain types of user profiles as  
6 explicitly indicated or what it would be for the collection of user profiles of those  
7 individuals within the present discussion groups (or if two individuals, the user profiles).

8 The large statistical information stored in SDI may be useful for the purpose of guiding  
9 discussions in a chat or forum context. Individuals may use the system in order to guide the  
10 flow of their own thoughts (where it is tuned to their own personal profiles) or that of what a  
11 particular individual with which they are corresponding at that moment (or in this case  
12 ideally both of their profiles) or (to appeal to the overall discussion forum or chat room) the  
13 collection of user profiles in that discussion forum. Again, in a variation, the system may be  
14 tailored to give additional weight to the particular historical experiences of the individual or  
15 group, thus if implemented in this way to the individual, the system may act in a similar  
16 capacity to a remembrance agent.

### 18 **1.3.11 Example: Vacation Package System**

19  
20 A vacation package organizer decides to begin a large-scale marketing campaign to target  
21 those people who would be the most interested in joining a new Caribbean Cruise. Although  
22 the vendor has a database of current customers, it is interested both in increasing the number  
23 and suitability of its potential leads.

24 Interfacing with the secure data interchange with which it is a member, the organizer  
25 identifies several possible sources of supplemental data: a LEIA-based travel discussion  
26 group, an on-line bookstore, and a Caribbean restaurant. These are found both by browsing  
27 through the interchange's internal list of members, and by using SDI-based data analysis  
28 tools, used within the interchange to automatically identify entities sharing common  
29 characteristics.

30 The package organizer then contacts each of these entities through the interchange, and  
31 negotiates different data-sharing deals: the travel discussion group is willing to exchange









1 by defendant , it such a equal mixture of empathizers from both sides constitute a  
2 virtual jury.

### 3 **1.3.14 Example: Group Therapy**

4  
5 There is a very useful and appropriate application for SDI to the automatic matching of  
6 individuals for purposes of group therapy. Individuals with commonality across many  
7 criteria can be selected, and multiple long-term groups adopted. In an application of this  
8 system, it may be possible to broadly extend the present scheme to the general public (where  
9 privacy is secured via the proxy server) and individuals may identify a group(s) which best  
10 fits their own unique emotional needs. It may also be useful to archive the sessions, index  
11 enabling the sessions and segments thereof to be searchable by keyword, e.g., via speech to  
12 text techniques and/or browsable by topical segment (which may be automatically  
13 segmented and labeled). The present system may also provide an ideal framework by which  
14 psychologists may identify patients of other psychologists that similar pseudonymous  
15 complements of psychological characteristics and symptoms, to allow targeted clinician  
16 interactions and robust and useful information from therapeutic approaches and/or drug  
17 treatment regimens.

18 In a variation, the present system methodology also enables a means by which much more  
19 specialized group session topics may be created which focus upon a very specific type of  
20 disorder, conflict or aspect of the patient's psyche. In this way, it may be possible for  
21 "identical" patients to, as group, focus upon different aspects of their ideal therapeutic  
22 regimen which collectively create a comprehensive customized treatment program. It is  
23 worth noting that the specialization possibly associated with particular focused sessions (and  
24 even types of individuals) provides a framework by which clinicians can become extremely  
25 specialized and expert within certain specific sub-domains of the field.

26  
27 It is certainly possible to adapt the above described technique for "smart interface" to  
28 provide useful ideas and suggestions for patients engaging in either psychotherapy or group  
29 therapy thus leveraging the information and personal experiences within the therapeutic  
30 processes of the collection of patients precoding them within similar topical sessions which  
31 shared very similar psychological and pathological patent profiles.

It is certainly reasonable and appropriate to adopt the present application framework to more several fields of medicine. For example, enabling physicians to identify other physicians have ----- had patients with the identical pathological profile and medical history. This enables them to ----- current medical history and clinical data, insights, observations, etc., in order for the physician to ----- the present clinical based on the expert advice and collaborative feedback from the other physicians possessing very similar experiences. Likewise patients could instantly access the physician or physicians which have or had experienced the most similar clinical situations that and the present situation of the patient in order to seek a second opinion. Physicians sharing unique clinical experiences (particularly clinically or scientifically in----- or series/problematic could use the ----- ----- scheme in order to ----- ----- companies, etc.. Finally one could also apply the -----  
-----  
-----.

### 1.3.15 Example: A Personalized Educational Portal

Another interesting application is to develop and deploy a personalized on-line informational portal containing everything from helpful links, illustrative content, text book information, quiz questions etc. The Personalized Educational Portal (PEP) typically is designed specifically for a particular class and includes several key features which are optimally applied as part of a comprehensive intelligent educational system. These features include the following:

- The ability to dynamically customize all forms of relevant information from the educational portal. The criteria for this customization however is based not upon the preferences of the user (unless for example the user receives credit for studies or projects or research on topics which s/he may select, rather it is based upon a predicted profile of the user reflecting his/her strengths and in understanding the relevant content. In particular the techniques of the issued patent number 6,029,195 [WHAT IS THIS PATENT NAME] describes a variation of user profiling in which

users are able to achieve a proficiency profile within certain domains of informational content where these informational domain(s) are determined in accordance with user's ability to answer a certain question(s) intelligently, discuss the answer to a certain question or about a certain topic or provide a useful reference or URL based upon the level of satisfaction of the requestor. This technique in itself could be usefully applied within the present application framework. For example, users may be students, and "experts" fielding questions could be other students (like tutors of sorts) and the payment they receive may be monetary compensation or even school, credit where other student's satisfaction ratings both qualify them for future opportunities to submit future responses within that particular knowledge domain as well as means of verification, and measurements of his/her proficiency over that particular material.

- Using the techniques of the parent patent application it is possible to also customize content delivery (including quantity, depth, and difficulty level) which reflects the user's strengths /weaknesses within the various relevant knowledge domains. It is possible to even statistically correlate the user's correct /incorrect responses to certain questions or types thereof as correlated with each other by common terms in the question answer part concurring in similar textual segments. These questions may be found in on-line exams or quiz questions associated with the actual content). It is possible to statistically correlate which content tends to best remedy certain deficiencies (as determined by incorrect responses to certain questions) by the user importance in being able to readily correctly answer those types of questions following reading the pages.

It may also be useful to apply the above technique towards determining which exam questions or combination thereof most commonly are revealing (if the student answers them correctly) of a student's proficiency within a particular small yet well definable sub-domain of knowledge. i.e. in which statistics show that If the student answers that sub-group correctly, s/he will likely answer the others in that sub domain correctly as well.

A sub-set of each of these exemplary questions from each knowledge sub-domain may be allocated as questions provided at the end of each relevant section of reading content (which may be on-line, off-line or only the responses may be entered on-line, exclusively

1 and if a wrong response is submitted, the system may recommend reading certain content  
2 which has (as above suggested) been statistically demonstrated to improve the student's  
3 proficiency in that sub-domain. It may also be possible to gauge in advance the students'  
4 predicted level and rate of advancement by knowledge domain via on-line psychological  
5 testing. Statistical correlations between these test questions and the student's learning  
6 abilities (by knowledge domain) can be usefully developed.

7 The system could even be constructed hierarchically by graduation of skill levels i.e. initial  
8 mix of questions of varying difficulty levels in order to gauge what level of content to  
9 recommend. The questions which follow are one level higher as is the next set of  
10 recommended content unless improvement is not observed.

11 It may be possible to a decision tree which automatically select which questions most  
12 effectively test the students, command of the materials (in this way a shorter exam may  
13 actually be better than for longer exams those questions which are more exemplary of such  
14 knowledge may carry greater weight). We could also effectively try to create a summary  
15 e.g. by topics /headings or even using text summarization techniques submit the areas of  
16 deficiency in one student to another student who is particularly proficient in that area for  
17 virtual tutorial. For this the tutor can receive monetary and/or also scholastic credit. The  
18 ideal scenario is actually creating virtual study sessions in which students are matched  
19 together which have the most different complement of proficiencies/deficiencies as possible  
20 and where these areas of potential mutual exchange benefit are maximized and are revealed  
21 to both students initially. Other forms of virtual study groups could be achieved by  
22 grouping students by similar proficiency areas and levels and creating a "virtual tour" with  
23 questions and content presented synchronously to a group of correspondence enabled  
24 students. In this way, if virtual tutoring or study groups occurs on-line exclusively, it may  
25 be performed pseudonymously as desired. We can even take a defined study group and  
26 assign it to group projects which can (again) relate to content to which the group is  
27 (collectively) deficient. (as group oriented problem solving has been proven to be extremely  
28 effective in remedying such deficiencies.

29 In a very novel application, it may be possible to even statistically using historical data to  
30 analyze the relative proficiency profiles of students by knowledge domain (perhaps) even  
31 further by, learning ability in those domains by teacher. Thus a virtual class could be

1 constructed using live streaming video which are truly customized to each student's  
2 particular needs and abilities, i.e. the idea would be before the semester to identify which  
3 teacher(s) were most successful in eliciting the highest levels of achievement among  
4 students which have an identical profile to that of the student (such statistics used in this  
5 profile may include but are not limited to the proficiencies / grades of the student in all other  
6 previous classes, psychological testing and/or (selective) knowledge domain proficiency  
7 questions (as above described).

8

9 Certainly depending upon such investor's personal preference, it is reasonable for each to  
10 agree within their personal data disclosure policies to disclose to SDI data regarding the  
11 types of investments, amounts and under what conditions (including those conditions stated  
12 by other investors such as the example above). The conditions for this disclosure, however,  
13 would be that no individual data be disclosed to another investor, i.e., that it be used only for  
14 SDI to be able to reveal aggregate statistics and predictions how the investment community,  
15 in aggregate, interacts with investment opportunities, and under what conditions, and in  
16 exchange for this disclosure, each investor earns the right to access this aggregate data. The  
17 individual investor or SDI, acting on behalf of that investor, may then suggest the best terms  
18 and conditions with which to negotiate with other investors, in order to optimally achieve  
19 their own desired objectives, part of the condition statement to SDI could also include the  
20 amount of money to invest in each synergistic investment opportunity which SDI typically  
21 would suggest to the investor based upon the relative distribution of its presently invested  
22 funds, size of the available non-invested funds as well as the relative risk/benefit proposition  
23 of the company to that investor (as estimated by SDI's use of company data as disclosed to  
24 that investor). SDI can even recommend or act as a proxy, in negotiating with the individual  
25 companies (who themselves could use an SDI negotiating agent representing their own  
26 interests). E.g., SDI, acting for the investor, could leverage considerable information about  
27 not only, the present investment opportunity, but literally all other investment opportunities  
28 which had been submitted to SDI for consideration by SDI affiliated investors. Based upon  
29 the total benefit which the investor stands to gain which is based, in turn, upon the SDI  
30 recommended comprehensive investment strategy) and the relative downside (relative risks





1 accept from the collective group of investors. Both sets of information (typically  
2 recommended by the SDI representation each associated party) are disclosed to the main  
3 (party neutral) SDI entity. If overlap exists in the range of terms or the ranges are close to  
4 one another, the introduction and negotiation between the parties becomes initiated (through  
5 the company(s) are not made aware of whether the overlap exists or is merely "within the  
6 general proximity" of accept-----  
7 -----introductions on the interest of the parties (though most markedly the  
8 companies)  
9

### 10 **1.3.16 Example: A "Group-Think" Ideas Market**

11  
12 In the consumer driven market place there is a significant unrecognized opportunity in  
13 being able to harness the intellectual capital of the consumer market. Historically, (in a  
14 non-networked environment) leveraging of such knowledge has been impractical and  
15 virtually infeasible due to the inability to identify and measure this resource of human  
16 skills (and thus match this skill with in the appropriate matching problem)with the  
17 particular types of problems which companies may be in need of new ideas and solutions  
18 as well as the lack of a trusted intermediary which can validate on an impartial including  
19 basis the authentic contributions of the provider of the information (as not having been  
20 claimed by the recipient after the fact as having been previously conceived before). Of  
21 course, there are remaining obstacles such as the fact that important problems are  
22 sometimes maintained confidential within a company (not even disclosed to some its own  
23 internal staff).  
24

25 With the emerging of extranets a certain amount of inter-organizational knowledge  
26 leveraging is becoming more readily achievable. A secondary advantage over an extranet  
27 is because of its ability to act as a trusted intermediary. It can validate the fact that that  
28 the individual is only dealing with one commercial entity and not its competitors (at least  
29 formally). Thus unlike extranet based knowledge leveraging it is conceivable a trust  
30 relationship enabling a certain amount of privileged information disclosure is at least  
31 theoretically possible.



(as similarly) may enable the automated or semi-automated selection of disclosed segments of users (containing particular attributes) which are likely to have relevant properties. The user selection process may be iterative based upon responses from users as they are selected. We can demonstrate this value experimentally, for example we can offer a vendor a free-trial and present personalized information/advertisements to one group of SDI users (both on the site and if desired also across the network), and regular advertisements etc. to another group without the aid of SDI. The increase in vendor revenue can be estimated from client-level monitoring of the change in purchase volume achieved with well-focused solicitations **on the vendor's own business**. Other prospective ads and syndicated purchasables could likewise be virtually overlaid on the vendor's site. Thus degree of increased click through rates, transaction rates and syndication revenues could be tabulated for that vendor in advance of his/her subscribing (or even being solicited) to SDI. Portals are also potential beneficiaries of this service in light of their need to better target their partner's ads.

The ability to the tailor a targeted marketing strategy to users knowing these correlations could be extremely beneficial in both on-line as well as off-line commercial environments. In the process of selecting the target user profiles significant value may be achieved from data collected from that vendor's own on-line presence, and a portion can be sold to other similar vendors; as discussed in the main description of SDI, this now forms a cached query which can be used by other vendors.

We might also demonstrate value to vendors with SDI by monitoring the performance of vendors with SDI technology, and providing metrics for new vendors to allow them to select suitable models of user targeting and personalization. Vendors that subscribe to SDI (the vendor centric version) can provide more attractive offers/products to users, based on information about the wider activities/interests of a user, on other vendor pages, and in the physical world (of course, only to the extent that this information is authorized by the user). Vendors can use information in the central SDI data warehouse for users' cross vendor and within vendor browsing and purchasing habits, and also with respect to profiling information about a user.

*One key application is first-time personalization*, so that information and products can be targeted to a user when he/she first hits a web page, based on profile information that the

user is willing to release. Vendors and users may also sell this information to other vendors. For example, information that a user likes a particular type of music is very valuable for vendors that sell content-based products, for example books, and CDs. Vendors can personalize their service, for example with collaborative-filtering based recommender systems.

### 1.3.18 Example: An Efficient Product Delivery System

The parent issued patent [FILL IN THE PATENT NO.] describes an application of collaborative filtering to the *strategic optimization of a vendor's business*, for example to allow a vendor to select an optimal location for an inventory warehouse based on projected consumer purchasing patterns; using the aggregate purchase history of users at that site compared with the other purchase selections at other sites for similar users. The model can also be used to predict demand for new items, and optimal locations for inventory given warehouse locations.

Now, within SDI the vendor can have control over two things:

- 1) the vendor has access to profile information about users
- 2) the vendor has a method to personalize and annotate on a dynamic basis the presentation of products to a user.

In common with the *smart caching* application of SDI to making communication networks more efficient, we can suggest an application to a product delivery system. At any moment a seller can have a good idea of the products that are available, and even the products that are in trucks etc., close to a prospective customer. This information allows the *vendor to selectively present items that are close to the user, with labels such as “this is within 30minutes of your door, NOW”*. Messages like this allow a vendor to manage its delivery process, because predictions can be made about likely purchase patterns, and then stock that is brought close to user locations can be explicitly advertised to those users.

Essentially, we attempt to predict transaction volume on a per item basis, and then positioning geographically physical inventory storage facility locations where inventory can be stocked so that items which a user is predicted to request are already located within the immediate physical proximity of that user.

1 As the trend to disintermediation continues (whereby we have informational middlemen, but  
2 not middlemen that physically hold products) we can allow vendors to notify portal  
3 intermediaries of product location within SDI, so that the portals can then pass this  
4 information onto customers.

5 The key role of the SDI system in this application is to allow users to receive personalized  
6 recommendations of products, and also product locations, based on their profile  
7 information—but without the vendors receiving that profile information. In addition, the  
8 shared information about user profiles in the data warehouse allows vendors to build good  
9 models for inventory location. Finally, SDI can act as a trusted intermediary between  
10 different competing vendors: for example suppose Amazon users bookstore A and  
11 bookstore B to provide books, and A and B do not wish to tell each other where their  
12 products are. The stores tell SDI where the books are, and SDI presents that information to  
13 users on a per-book basis, as books are browsed within the purchasing system.

14 There are numerous useful applications to improving quality, speed and cost of delivery to a  
15 user. In one example, it may be possible to provide same day delivery for on-line purchases.  
16 We could also restock a truck on the basis of what a local population of users are likely to  
17 purchase; with items presented to a user along with a particular anticipated delivery period,  
18 with the Global Positioning System (GPS) or more LEIA [INSERT THE PATENT NO.]  
19 on the vehicle providing up-to-date information about a truck's location and anticipated  
20 delivery time.

21 We might also suggest that a user can physically travel to some location close to his/her  
22 base, based on information within SDI about the user's travel patterns, etc. Purchasable  
23 items can then be located in locations that are convenient for users, or shipped dynamically  
24 to those locations (i.e. intermediate warehousing systems). Users can be provided with real-  
25 time directions to the location of such a pick-up point via personal digital assistants (PDA)  
26 and navigational devices, etc.

#### 27 1.4 Push-based Advertising/Solicitation

28 Another form of information that a user can submit to the central SDI database via his/her  
29 client is the advertising acceptance functions, which state the explicit preferences of a user  
30 for adverts, and required payments. The level of compensation that a user requires to receive  
31 an advertisement will depend on the relevance of an advert.

1 An application of SDI is that vendors can use the centrally stored information to request that  
2 as a result of a query the query-execution module in the central SDI data warehouse sends  
3 advertising solicitations to appropriate users, i.e. users that are willing to receive adverts as  
4 specified within their acceptance functions, and willing to release profile information to  
5 enable useful advertising as specified in their price rules for regular data.

6 Figure 15 illustrates the system for push-based advertising. A vendor requests the right to  
7 advertise to users by sending a special type of query to the central SDI data base, where the  
8 query states the vendors preferences for user profiles and requests that adverts be sent to  
9 users. The vendor never receives the contact information for user's directly. The server also  
10 receives payments from vendors and credits users' accounts.

11 Suppose Vendor V wants to advertise a new product to relevant users. Vendor V can use  
12 the SDI database to select users with a high predicted hit rate for the advert, and then  
13 compute the value of a bid that it is prepared to pay users for the right to provide them with  
14 its advert. The vendor makes a request for bids to appropriate users by formulating this as a  
15 query with an associated action, and submits to the central SDI data warehouse where it is  
16 handled by the query-execution module. Part of the query is a bid function, indicating the  
17 maximum amount that a vendor will pay to display an advert.

18 The query-execution module evaluates the request with respect to the profiles of each user,  
19 the rules that agents submit for allowing access to profile information, and the advert  
20 acceptance functions. If successful the vendor pays the minimum value necessary to have its  
21 ad accepted by the user, again simulating a Vickrey auction as for the dynamic competition  
22 for banner ads described earlier.

23 For example, suppose a vendor has an advert **Ad** to push to users, and is willing to pay up to  
24 \$1 to place the advert with users that have a predicted hit rate of greater than 30% on the  
25 advert. The vendor can formulate a query that will first identify users with a predicted hit  
26 rate that is greater than 30%, and then determine which users will accept the adverts at the  
27 bid price. If successful, the adverts are pushed to the users via the SDI central data  
28 warehouse, and the payment required by each user is submitted, less than \$1 in all cases. In  
29 a simple variation, the vendor might also state a fixed budget, so that it does not send  
30 advertisements to more users than it can afford.

31





- 1 (i) The vendor assesses the value of the information present in the secure data  
2 interchange. This computation is performed securely either by revealing randomized  
3 aggregates to the vendor to enable its own local analysis, or by allowing the vendor to  
4 check data and algorithms into the secure data interchange site for analysis.
- 5 (ii) The vendor selects criteria for mailing unsolicited advertisements, and agrees on a  
6 pricing model. In this case per-impression pricing is the most obvious pricing model,  
7 as it is difficult to monitor when a user responds to unsolicited mail per-transaction  
8 pricing is difficult. The user could be motivated to do this should the Secure Data  
9 Interchange promise future returns for recording a successful solicitation with the  
10 database.
- 11 (iii) Either the data list is released to the vendor for its use, if this is within the selling  
12 vendor's data policy, or the data interchange sends mailings on behalf of the  
13 purchasing vendor.

## 14 1.5 Community Dollars and Business Models

15  
16 The primary objective of the iamworthit model is to create a market for information  
17 about buyers, with agents that submit profile and other personal information to the  
18 database able to collect payment in return for queries executed by vendors. In an  
19 important extension, we can allow this payment to be made in terms of *community*  
20 *dollars*, which can only be spent with particular vendors. Community dollars allow a user  
21 that provides profile information to the system of SDI to receive payments that are  
22 *dedicated for a particular type of purchase*, these payments are called "community  
23 dollars". One central example is to allow a vendor that signs a user into the SDI system to  
24 be able to lock a proportion of payments accrued by that user to the vendor's own  
25 product/service domain. This is important, because vendors now have incentives to bring  
26 new users into the system. A vendor can offer a user community dollars on its (and its  
27 affiliates) web site, in exchange for receiving user data via the SDI data exchange.

28  
29 We allow vendors to pay in "community dollars" for adverts, dollars that can only be spent  
30 at that vendor (with the host site of the advert receiving a share of the profits). This provides  
31 vendors with the ability to gain long-term customers. Furthermore, so long as the buyer

1 agrees to receive advertising from his/her iamworthit subscription offer, community dollars  
2 can be replenished at the rate at which advertisers are willing to pay for impressions. This  
3 provides buyers with an incentive to spend at the vendor's site, because the vendor can  
4 monitor (pseudonymously) the buyer's that are sensitive to discounts and other special  
5 offers (that are delivered as community dollars). With community dollars a vendor can  
6 compensate buyers for information that they access, but tie that information to certain  
7 vendors (e.g. the vendors that first signed the user into the system). The system has the  
8 following useful properties:

9 (1) buyers are incentivized to provide information that allows vendors to push relevant  
10 advertisements/products;

11 (2) buyers will also be more likely to make purchases at a site or a coalition of sites for  
12 which they can receive discounts via community dollars;

13 (3) providing buyers with community dollars will increase the number of hits to a site.

14 We allow buyers to receive compensation for providing personal data to vendors,  
15 information that has value to vendors because it allows information to be focused (for  
16 example relevant ads can be displayed to a buyer, based on his/her profile). The system of  
17 iamworthit credits buyers for information, and provides buyers with direct incentives to  
18 reveal profile information to vendors.

19 A vendor can sign up with iamworthit.com and agree to provide only the most restrictive  
20 type of community dollars, that can be spent at that vendors site. Community dollars are the  
21 currency that vendors provide in return for the right to provide focused information to  
22 buyers. Dollars can be general (e.g. for a network of vendors), or very tightly focused (e.g.  
23 for a particular product, at a particular time). The buyer-centric infomediary acts as a broker,  
24 matching buyers and vendors. Another key role of the infomediary (e.g. the portal) is to  
25 protect the buyer from information saturation by controlling the flow of solicitations. (i.e.  
26 restrict the number of ads. that a buyer sees)

27 Community dollars extend the methods in Secure Data Interchange that allow a user agent  
28 to sell controlled access to information to other agents, introducing in addition a method  
29 which can control the ability of a user to spend received payments. For example, a vendor  
30 might be happy to pay a user for the right to display a targeted advertisement if the payment  
31 represents a discount that can only be used against products offered by that vendor.



1 for offers which are delivered via digital coupons and/or as “straight value” which could be  
2 converted directly to purchases thus are equivalent to real dollars at the point of transaction.  
3 The community dollars can be “credits” that can be redeemed as real cash, credits towards  
4 discounts, and can be spent across a suite of sites, or limited to one site. The co-pending  
5 patent application entitled “System for the Automatic Determination of Customized Prices  
6 and Promotions” [INSERT THE PATENT NO.] describes a comprehensive scheme which  
7 may be implemented in either on-line or off-line commerce environments. The system  
8 enables vendors to deliver a digital message in the form of a promise to a buyer (typically on  
9 encrypted form for purposes of targeting a buyer specifically). This promise is typically a  
10 discount for a product, set of products (or all products in stock) or may even include  
11 entitlement to special privileges for that buyer, thus it is termed a “digital coupon”. The  
12 community dollars can represent special discounts for a buyer.  
13 The buyer receives a financial incentive for receiving well-targeted solicitations, while  
14 preserving buyer privacy within the SDI system. The vendors support the community  
15 dollars through advertising revenues and increased sales volume. We can also provide the  
16 vendor through which the buyer first subscribes a special “first screen” right that allows the  
17 vendor to provide a buyer with his/her first impression as soon as s/he logs on.  
18 Of course many variations of the community dollars scheme are conceivable and the current  
19 description is in no way intended to limit the scope of the claimed invention. For example,  
20 the ad revenues generated may instead be apportioned between direct payment which the  
21 buyer receives, in real cash, community dollars which the vendor (or ISP) credits the buyer  
22 and/or direct payment which the vendor (or ISP) receives as well as iaworthit or any  
23 combination of the above such as exclusively direct payment which is apportioned between  
24 the buyer and the vendor (in the absence of the community dollars scheme).  
25 In one variation all community dollars collected by a buyer must be spent back at the vendor  
26 site at which they originally subscribed (and also the site that hosts the adverts of other  
27 vendors). A buyer can spend the dollars with any vendors that are site partners of the  
28 original site. This provides the vendor an incentive to accept and promote the community  
29 dollars concept.  
30 The value of providing a buyer with targeted solicitations is estimated at approximately  
31 \$300 to \$500 per year (based upon \$120 per 1000 targeted impressions at approximately 25





which is possible via the one-way identity look-up functionality of the central SDI data warehouse. The eBank can query the look-up table, and check that both PIDs correspond to the same UUID.

For example, an amount of dollars below a certain threshold can be programmed to have no value, but additional dollars can have increasing marginal value. This provides a non-linear incentive for a user to stay with the same vendor in a market with low switching costs, or a non-linear incentive for a user to receive adverts from the same advertising agent. The advantages to vendors and advertisers are repeat purchases, consumer lock-in, and also improved profiling from interacting with the same agent over extended periods of time.

The functionality is embedded into the community dollars. The framework expands the idea presented in Chaum [Chaum 85; Chaum 92] where the public key of a public key /private key pair indicates the dollar value of a community dollar. In this case, we allow any number of public-key/private-key pairs, and use the public key as an index into a table maintained within the eBank to provide additional criteria about how the cash can be redeemed, for example it has more value if used in combination with other community dollars, it decays over time from date YY/YY/YY, etc... We do not intend to limit the variations possible, but provide this as a framework for ways in which constraints and conditions on community dollars can be extended. For example, community dollars can also be programmed to lose value over time. This can increase the number of repeat purchases

In the preferred implementation we use an “electronic cash” infrastructure for the community dollar system. A buyer’s SDI-enabled client-level proxy stores dollars that the buyer receives securely. Dollars are anonymous and non-traceable, so that the buyer can maintain a single “bank” of dollars, and aggregate dollars collected across pseudonyms for a single purchase, so long as the purchase satisfies the constraints on the dollars. Each dollar is created using Chaum’s blinded signature technique, and also signed with the conditions on its use.

This scheme allows vendors to monitor the offers that buyers respond to, because when a buyer presents a community dollar—the dollar can be validated to indicate the type of discount that it is, even if the identity of the dollar (i.e. the serial number) is untraceable. SDI provides vendors with guarantees that buyers have once-in-a-lifetime pseudonyms, so redeeming a voucher of a particular type that is redeemable only at vendor  $V$  and was issued







1 transactions. There is considerable economic pressure on web tool providers to provide  
2 competitive solutions which are not only rapidly deployable but also extremely robust  
3 (typically bundling some form of personalization technology), thus as a result of  
4 economic pressure to drive advertising and e-commerce transactions through the site  
5 (often through the use of a revenue sharing model which can reduce upfront costs to the  
6 vendor substantially). iamworthit through its customer loyalty enhancement capability  
7 provided by community dollars can significantly improve the effectiveness of  
8 personalization at the site (perhaps more significantly than cookies which often are  
9 blocked and all typically not implemented so as to be recognizable across the collection  
10 of sites which the tool provider supplies or the web host services).

11 It is believed that the revenue sharing model may be structured to cover both the cost of  
12 the tools, web development services as well as web hosting (thus unless he himself  
13 offers a complete turn key suite of solutions and services a tool provider which  
14 integrated the community dollars concept could be in an extremely competitively  
15 advantageous position in attracting strategic relationships with Web developers and Web  
16 hosts. The tool provider may for example sell an ad server module or could  
17 "automatically subscribe" the user with an opt-out option. The very fact that all of the  
18 tool , Web development or hosting provider's sites are exempt from ad blocking/ad  
19 replacement technology (barring iamworthit competitors) would be a significant  
20 motivating factor in incentivizing sites to utilize their services (e.g. through promotional  
21 programs). An emerging large market for Web development is providing these services  
22 to individual end-users. Completely free services in this regard would likely mobilize  
23 what activity in what is now a dormant, albeit enormous commercial market. For  
24 example, tools for developing elaborate Web-wide community portals could have  
25 customized community links which are matched to the user's personal profile. Visitors  
26 to the site (subscribed to iamworthit) could experience an additional layer of  
27 personalization based upon their own user profiles.

28 Web hosts may also use an additional feature (optionally to significantly drive increased  
29 traffic through potentially all sites on their network. This optional feature is a set of  
30 links (e.g. along the side of the user's screen). These links are used to point the user to  
31 other relevant pages issuing aggregate site usage statistics of their visitors (as described

1 in the parent patent application) iamworthit, user referral links may additionally  
2 personalized based upon their user profiles.

### 3 **1.5.4 Vendor coalitions**

4 Vendors may choose to form coalitions, to allow buyers to spend community dollars at  
5 any "partner" site. Vendors that have similar buyer bases can be automatically identified  
6 using collaborative filtering. (i.e. determining similarity with the present vendor, from the  
7 aggregate vendor preferences of a given vendor's subscribers). Also, these resulting  
8 metrics could incorporate predicted online spending by each buyer at each site. This  
9 could help to narrow the selection of sites the vendor wishes to partner with and/or the  
10 selection of these partner sites could be determined and presented to the buyer to even  
11 further narrow the selection for each buyer. All vendors in a coalition advertise, and  
12 provide cross-links and up-links to other vendors.

13 The coalition model is good for buyers that are more likely to find products that they  
14 want. Vendors can share the risk of advertising, since dollars provided to one buyer by a  
15 particular vendor can be redeemed at another vendor. Advertising and community dollars  
16 increases sales volume at all vendors in the coalition. Furthermore, studies on on-line  
17 buying behavior suggest that on-line shoppers tend to make purchases across a variety of  
18 categories of e-commerce products, thus it is likely that they would also prefer the  
19 freedom and flexibility which is associated and can be provided in this way only by  
20 deploying a multi-site community dollars scheme. Iamworthit usage statistics are very  
21 effective in identifying prospective vendor coalitions which consist of complementary  
22 (non-competitive) vendors. Alternatively, personalized coalitions which can be created in  
23 ad-hoc fashion for each buyer can further serve the buyer's interests by increasing the  
24 freedom and flexibility which the buyer often strongly desires. Though practically this  
25 can only occur for those (perhaps smaller less commercially significant) vendors which  
26 have not been a major coalition consisting of vendors which many buyers tend to use  
27 which collectively serve a substantial percentage of the overall user base.

28 Thus if the vendors have not established firm partnerships with other vendors, we can  
29 even allow vendors to form dynamic and virtual coalitions within SDI, with a potentially  
30 unique coalition of vendors for each buyer. The coalition may consist of an optimal pool

1 of vendors, as determined by SDI collaborative filtering techniques. The goal in this  
2 model is to provide buyers with a particular “brand” of community dollars.

3 • Multi-vendor Community Dollar Portal

4 We can allow each vendor to retain an exclusive right to advertise to each buyer; and  
5 also develop a portal for the coalition—that gives advertising prominence to coalition  
6 members. Portals will be expected to aggressively promote community dollars.

7 Buyers that collect community dollars become loyal return visitors to the portal and  
8 its associated vendors. In the case the vendors do not generate the same value we can  
9 provide community dollars in proportion to the value that a vendor contributes to a  
10 coalition.

11 We can also provide targeted advertisements for the vendors at the portal, using the  
12 user profile to focus ads. The categories and links at a portal (that might include a  
13 search engine) can be re-prioritized (highlighted and/or re-ranked) in accordance with  
14 the buyer’s preferences (as described above), and to favor subscribing vendors. In  
15 commercial practice, the motivation for vendors to become coalition members is  
16 largely driven by customer partnerships with providers of e-commerce tools and  
17 solutions. There are many emerging trends by which these intermediary tool  
18 providers could conceivably integrate community dollars. Consider loyalty points  
19 (e.g. [www.mypoints.com](http://www.mypoints.com)) or loyalty bonuses programs (e.g. [www.clickrewards.com](http://www.clickrewards.com)).

20 These points or bonuses could be substantially subsidized by the advertising and  
21 deeper benefits passed onto the consumer. Other simpler technology which would  
22 compel these vendors to cooperate include shopping basket technology, the emerging  
23 standard, ECML, common Web tool and/or tool development solutions, common  
24 hosting solutions, common ad delivery systems.

25 Vendors pay the portal site to advertise, and the portal provides community dollars to  
26 buyers in return for privacy-protected profile information. This model does not  
27 provide incentives for the portal to provide well-targeted adverts, because there is no  
28 direct link between a portal’s revenue stream and the vendors’ sales volumes.

29 A portal with community dollars that can only be spent under a single pseudonym at  
30 its partner sites also provides an incentive to buyers to interact under a single  
31 pseudonym—which in turn allows a portal to profile buyers across its complete



types of commercial partners) in order to dedicate a certain percentage of the community dollars (e.g. thirty percent or approximately one hundred fifty dollars per customer) which could only be redeemed at that multi-store retail site (and/or the value of these dollars could be worth more at the retail site). In addition, in this model, the independent advertising initiative of iamworthit would be geared towards community dollar credit of that retail site. It should be noted that, because if other outside competition occurs to the basic iamworthit scheme to a substantial degree there will not be a compelling incentive for buyers to adopt a more restricted form of value (as retail credits at a particular site), versus accepting the credit from a competitor in the form of cash. Thus this model could provide a viable means for attaining a leading position in one or more on-line retail markets if this competition does not substantially exist.

- **Marketing Network.** Iamworthit sites which offer a community dollars promotion could, upon the buyers subscribing to iamworthit, additionally offer the buyer with a down-loadable client based software which provides a small promotion in conjunction with a link to iamworthit. Each time a recipient of the email subscribes to iamworthit, a percentage of the value of that customer is credited back to the buyer in the form of community dollars. Each subscriber resulting from the current subscriber's email (though reduced) provides an additional credit to the original subscriber in accordance with the marketing network business model. If the site originally delivering the promotion is not an e-commerce site, a percentage of advertising revenues resulting from the subscriber (and potentially all resulting subscribers) could be used. It could be applied in the form of iamworthit advertising (or exchanged) for advertising in an ad server.
- **Free Web hosting.** A portion of community dollars may be allocated in a revenue sharing arrangement between the buyer and the web-site from which subscriptions to the service are provided, thus enabling Web hosts to offer free hosting services while also receiving full payment which is iamworthit advertising supported. Iamworthit could also easily use its vast statistical data collected from each site in

the hosting network in order to predict which sites tend to be most commonly linked to from the present site (and/or share “similar content”). This technique is described in the parent patent application by establishing these links users can better access relevant information . In one preferred variation, all non- customer’s sites are not included in the referral links. This level of targeting is likely to drive significant traffic through the Web hosts’ network (and even more so if iamworthit user profiles are transferred to the site upon visitation of an iamworthit subscriber as these similar cross-links can also be personalized to the user.

### 1.5.5 Transaction-based Revenue-sharing

In transaction-based revenue sharing, the only time that advertisers pay to provide an impression to a user is when a sale results, in this case the hosting site receives a cut of the final transaction price. The vendors provide buyers with community dollars directly. The dollars, which are stored at the portal site, allow buyer-spending to be tracked. This allows the portal to monitor when a sale occurs, not just a hit on a banner ad. With transaction-based revenue of this kind, personalization is critical. In this model the portal will give prominence to adverts from successful sites. A portal site may forgo payment from a vendor in exchange for the increased click-through from a strong network of community-dollar enabled vendors. Value is credited directly to buyers for future redemption at that particular vendor's site. The community dollars provided to a buyer can be restricted, such that a buyer can only redeem dollars if s/he maintains enough visits to the portal site.

Vendors can offer discounts on their own products directly, instead of providing the portal with money for advertising. The vendor only pays to the extent that its advertisements are well-targeted. The vendor could also request special ad priority. A vendor that presents advertisements to a buyer offers the buyer discounted promotional offers for products offered by partner vendors, in exchange for subscribing to iamworthit and receiving targeted impressions. These offers are in lieu of community dollars, and can be provided by partner vendors—maybe in exchange for a right to a number of ad deliveries for the vendor’s own advertising purposes.





1 implement this monitoring, because no other system knows a buyer's portfolio of  
2 pseudonyms. The buyer can present its digital credential when visiting a vendor's site.

3 A vendor may wish to provide loyalty dollar credit; for example, it would be possible for  
4 vendors to offer buyer's credits if the buyer is a 100% loyal customer i.e. that she/he did not  
5 (over a specified period) do his/her purchases at the site of any competitor. For example,  
6 certain types of high value customers could be given considerable value in the form of  
7 credits or discounts as a result of demonstrated vendor loyalty. The credential can be time-  
8 stamped, to prove loyalty. It does not reveal any information about the buyer's other  
9 pseudonyms to a vendor, because many pseudonyms will exist that have not made any  
10 purchases from a competitor. Upon accessing the vendor's site, this credential may be  
11 presented to the vendor. One criteria for the above benefits could be that the buyer may visit  
12 a competitor site, and engage in interactions; however s/he should not transact with that  
13 vendor.

#### 14 1.6 Alternative Business Models

- 15  
16 • Offer discounted or free services in return for the right to access profile  
17 information.

18  
19 Examples: Free or discounted retail products with "niche" partners in each category;  
20 Free dial-up ISP (as an independent ISP or a service to jointly promote free access  
21 with ISPs); Free Cable and ISP service, Free pay-per-view (note that viewing patterns  
22 and the associated content could provide additional valuable user profile  
23 information); Free phone service (e.g. advertise subscription service on screen phones  
24 or audio ads from pay phones); Free prepaid calling card; Free print media  
25 subscriptions (magazines, newspapers); Free book clubs; Offer any combination of  
26 the above with "deep discounts" for each (this can involve \$350 community dollars  
27 per buyer or it may simply involve certain purchasing limitations per customer).  
28 Each vertical niche partner in exchange gets exclusivity within their own respective  
29 niches to target advertise to those buyers (e.g. retailers); Free access to sporting  
30 events; Free credit for casinos; Free lottery tickets; Free charity donations;  
31 Discounted hotel lodging; Monetary credit to a credit or debit card (either an

iamworthit branded card or provided as a partnership with the card companies;  
Monetary credit to a diner's club; Free subscriptions plus credit to retail buyer's clubs  
(on-line or off-line); Credit or discounts for book clubs; Free musical concerts, or  
theater presentations, movies or access to arcade entertainment; Free access to  
amusement parts or theme parks; Free golf season passes; Free commission fees for  
stock trading; Free commission fees for travel booking (if implemented for on-line  
buyers would be less compelled to search for travel information on-line though go  
off-line to make their bookings).

- Allow the ISP to promote free Internet access through personalized advertising on [iamworthit](#).

The ad server can even recognize through the associated domain names, the buyers which are coming from a competitor ISP. So long as that ISP is not a partner of iamworthit, the associated buyer would be selectively targeted with an offer of this sort “free Internet access” by subscribing to iamworthit”. Smaller ISP’s would be particularly compelled by such offers to their direct target prospects. This is because they are operating on a “thin margin”. Furthermore, both they and their small regional counterparts would be particularly vulnerable to this type of advertising by regional competitors from the same geographical area, during specified period of months of initial usage of the service, the share of profit due iamworthit could instead be committed to purchase additional advertising for the Internet service provider (or the balance of this profit traded-out in the form of additional advertising through the ad server partner).

The ad server partner could further become an exclusive partner of iamworthit on the following commercial venture: Relationships as established with on-line merchants and other e-commerce sites. The vendor actively promote an offer to their customers through both off-line media (using a URL unique to that vendor) and on-line advertising through the ad delivery partner. The offer may say (as an example), "receive three hundred dollars in purchasing credit at Books a Million in

Alternatively, it would be possible to offer websites the ability to become Internet service providers where the interface to the ISP home page would essentially be heavily branded to that site or portal. Companies like GTE already offer a “Virtual ISP” service in which the content to the ISP home page is unique to the ISP while the network is provided by the virtual ISP service. This model would be particularly compelling for sites which are largely community oriented and have a potentially loyal customer base. Moreover, interestingly, many of these community sites are offering many of the services and capabilities that a full-blown ISP would offer from its home-page, e. g., a portal interface, links to high-quality content, chat/forums, e-commerce, commerce affiliate links, etc.

- **Free Community-based Content.**

Create premium content which is free to iamworthit subscribers, subsidized by revenue from profile information. Each iamworthit buyer would be granted free access privileges to the premium content on all sites which are part of the program.



- Allow matchmaking of buyers based on profiles across a network of iamworthit/community-dollar sites.

A menu of different forums and chats can be displayed on each iamworthit member site. (the underlying methodology for which is described in co-pending patent application “Virtual Community Service for System for Customized Electronic Identification of Desirable Objects”). In accordance with this specification, a variation of the service involves the process for identifying individuals who most closely match a given category or target object. For example in the context of the present implementation a category or content, merchandise or a purchasable being specially promoted may be the focal point of a discussion forum or chat room, which is automatically organized by the Virtual Community agent. Accordingly, a portal (or in accordance with the present trend) a site with which a portal interface is integrated utilize the present techniques for generating virtual communities for each category or sub-category of content on the portal or for direct access into a forum or chat room which was automatically created around that particular site (as the target object used as the matching criterion). As described, the user may navigate a hierarchical menu of virtual communities which may be constructed automatically according to the methods described which involves communities assigned to category, sub-category,

and association with corresponding sites. Ideally in this scheme the portal is actually a “virtual portal” which may be utilized in providing access to the communities across numerous sites (and/or ISP home-pages). Users may also be navigated (at the individual user level) which along with their pseudonymous user profile data is subject to their data release policies. In a variation of the above schemes, if there is geographical information which is associated and which is released in accordance with the above individuals and/pr communities (e.g., as may be occurring or scheduled to occur in physical space), LEIA may be employed as a primary (or additional) selection criteria for navigating the present information accordingly.

- Advertising in Exchange for Equity.

A potentially attractive optional form of value, which could be provided to iamworthit customers involves equity shares in companies which advertise to the buyer (in lieu of community dollar credit or cash). This scheme is an ideal application for iamworthit in as much as iamworthit customers can be highly targeted and because many Internet-based start-ups are highly niche community oriented (thus iamworthit customers who are interested in the sites can be efficiently identified and targeted). Moreover advertising is typically very expensive which in the absence of accurate targeting may be of questionable value. It should be noted, however, that because the primary objective is to both find viable prospects and to engender an element of loyalty (which the equity model does). This scheme would be the preferred approach to advertising for sites which do not sell on-line where community dollars would be the preferred loyalty engendering scheme. In order for this model to substantially provide its desired advantages of increased advertising exposure to fledgling web based companies, the iamworthit subscriber base would have to be quite substantial.

- Loyalty credits for off line retailers.

Deliver through the back of sales receipts, kiosks or direct mail or on-line substantial purchase credit to retailers (e.g. grocers') customers, using the aforementioned technique of utilizing a unique URL to identify the vendor and/or promotion from which an iamworthit subscriber originally accessed the iamworthit subscription site (thus identifying for both buyer and vendor the appropriate denomination and/or terms of community dollars issued to the buyer). In the preferred implementation, a loyalty card is used to identify the buyer thus enabling the community dollars value to be provided to the customer at check-out as straight credit or possibly an enhancement to loyalty credit. The buyer may also be identified via credit card or alternatively a voucher (or coupon) could be printed from the buyer's computer or from a kiosk which is typically situated near the entrance to the store and which could be activated upon insertion of a loyalty card credit card (or associated authorization code) and could also be used to disclose the buyer's community credit balance. A unique identifier for that voucher or coupon is provided and non-tamperability measures are provided such that the buyer's community dollars account can be appropriately debited upon redemption. Preferably, a pre-determined value is specified on each voucher (which could be predetermined by the service or the buyer) or alternatively, the total community dollars balance could be specified on the voucher along with the buyer's name/address and redeemable only upon presentation of valid buyer ID.

- Free ASP Services – Web-centric applications are becoming an increasing central part of the e-business infrastructure. iamworthit could potentially enable organizations to gain free access to these Web-centric applications in exchange for iamworthit enabling their employees and customers (e.g. which may use some of these applications as well).
- Credit to a User's Credit Card – Many credit cards are tailoring promotions to enhance not only acquisition of market share but also loyalty of its users. The user-centric SDI provides an enhanced (e.g. rebates for expenditures) platform which could provide more efficacious loyalty enhancement marketing strategies

for card issuers. For example, a credit card user could be given certain credit towards the card (e.g. via a promotional offer) for becoming an iambworthit subscriber (which is largely ad revenue supported). An additional novel application could involve providing an offer for an additional benefit which could be provided if the user performs all of his/her on-line transactions using the present credit and (i.e. wherever the card is accepted).

An iamworthit pseudonymous credit card [e.g. LMP 94] could be provided whereby users collect credit for advertising (e. g., on billing statements for the ad) which the user receives as well as the advertising delivered via iamworthit's on-line (Web or targeted email) as well as pseudonymous physical mail, telemarketing calls using user profile pseudonyms and one all time or pseudonymous phone numbers. This user profile data consists in part of off-line data and on-line data.

- **Free Personal Portal for Individuals.**

It is reasonable to offer individuals completely free, Web design/development and hosting services which are offered and mass marketed. At the time this patent was written, Web developers were offering such free services in exchange for revenue sharing for advertising product syndication or e-commerce offering on the individual's site. Iamworthit could conceivably be deployed in conjunction with this commercial model whereby either a portion of the user's own iamworthit profiles are used to subsidize (or subsidize in part) the portal and the developer and/or revenues shared from profits from iamworthit subscribers who subscribed from the portal are utilized also (or instead). As such, it is also reasonable for such a developer to provide links on affiliate sites (e.g. access an affiliate network) which offers free individual portals in exchange for subscribing to iamworthit (and perhaps agreeing to offer a similar iamworthit expectation from that portal as well from which the user can also share in a portion of the revenue stream. Such a program could even be structured in its revenue distribution to the subscribers as a multi-level marketing network. Users could even receive value for providing links to the vendors site either



both (e.g. as “cob----- shopping portal) ----- direct off-line encounters in the other users as prospective customers to the URL. (where the URL is unique for each individual.

- **Multi-level Marketing Network.**

This economic model for distribution of shared revenues from iamworthit may be implemented several ways. For example, as above suggested, the user's Web portal offering solicitations for iamworthit (as a standard offering or in combination with a free personal portal), could be implemented in a number of different ways, such as ;

a). The solicitation is presented upon the user's Web portal.    b). The solicitation is made through the traditional sales channel of direct one-o-one correspondence with friends or associates.    C). The solicitation is attached to any email messages of the user such solicitation provides a link to enable the user to subscribe.

### 1.6.1 Agent-Mediated Value Exchange in the Supply-Chain

The information in the SDI data warehouse can allow vendors to identify business relationships between other vendors, for example between vendors and their suppliers, and between vendors and their customers. With this information a vendor might attempt to exert pressure on another vendor by making direct offers and incentives to their suppliers and/or their customers. Of course, in all cases information is only available subject to price and disclosure policies, but with there being at least two parties in any deal no vendor can unilaterally prevent this kind of information from reaching the SDI enabled data market place.

Information may include the benefits/disincentives associated with particular actions, and help to identify actionable events. Agents might offer other vendors a share in some future profits that result from taking a particular action, so that little up-front collateral is required. The information in the data market can be used to estimate the value of this

1 share, for example, based on information about expected trade volumes from projected  
2 market share.

3  
4 In a variation, SDI may also assume the task of negotiating on behalf of each entity  
5 involved. This is useful when information is sensitive and/or confidential, SDI can allow  
6 negotiation without identification until a deal is struck. In this section we describe an  
7 application of SDI to an agent-based economic infrastructure in which value exchange is  
8 used to create "efficient supply chains and contracts".

9 It is critical within this system that agents can specify quite elaborate privacy policies, so  
10 that they can submit sensitive information to the SDI data warehouse but be sure that the  
11 information will not be revealed to competitors. One consequence is that SDI will compute  
12 efficient supply-chain structures based on the information provided by agents, but be able to  
13 report a good set of contracts *without providing complete rationale* (e.g. plans, strategies and  
14 or future technologies) for the basis of the identified solution. An agent may be asked to  
15 apply economic pressure within a supply chain without understanding why it will receive  
16 personal (and perhaps long-term) benefits. In another variation, if SDI is able to keep track  
17 of an entity's competitors and cooperating partners, then certain amounts of information can  
18 be carefully released to justify decisions.

19 SDI must retain a considerable amount of autonomy, to apply its knowledge to act in the  
20 best interest of these entities which it represents without being able to disclose the  
21 rationale for their contract recommendation. There is a very significant amount of trusted  
22 information regarding all of the commercial activities, internal operations, marketing ,  
23 strategic business and product development strategies, etc. which must be disclosed to  
24 SDI for an accurate assessment regarding the present and predicted commercial impact of  
25 new supply-chain contracts.

26  
27 Furthermore, we must prevent a vendor from exaggerating its value of certain outcomes,  
28 to achieve a good solution. Given this *incentive-compatibility problem*, and the potential  
29 problem of vendors entrusting so much extraordinarily delicate information to a single  
30 entity, we could instead allow SDI to assume many of the roles of a *consulting firm*. In  
31 this case SDI controls the reporting of information to SDI, and can verify its accuracy.





1 With interwoven supply-chains, for example, with the same supplier connected to  
2 multiple vendors, and vendors with competing suppliers, this method of passing value  
3 along the supply chain becomes more tricky. When a consumer applies economic  
4 pressure on a vendor that it deals with in the supply chain, that vendor may inform other  
5 suppliers of the pressure and block service to the user. A bartering system can provide  
6 additional benefits, to allow vendors to exchange bonuses, discounts, etc. with favors  
7 offered to vendors in other value chains. In general it is quite likely that a favor will not  
8 substantially benefit a user in another supply chain. We can also include a method to  
9 prevent a vendor that does not itself offer favors in the supply chain from taking  
10 advantage of favors offered by other vendors. This can support cooperation  
11 within the supply chain, and allow effective competition with other supply chains.

12  
13 Each vendor in the supply chain can benefit in two ways:

- 14  
15 1) They receive a “commission” on the transactions in which they convey the favor  
16 to their customer (which again may be “upside” in the ultimate beneficiary of the  
17 favor and/or product/service loyalty credit with their own supplier).
- 18  
19 2) (Most importantly), they receive customer loyalty advantages via the community  
20 credit they provide such customers.

21  
22 For reasons of this latter benefit they are likely to compete with other competitive  
23 vendors for the privilege of exclusively delivering the incentive to evoke the favor to the  
24 next level up the supply chain. In this way each vendor below the supplier requesting  
25 stands to gain via customer loyalty benefits by competing with another supply chain.

26  
27 A consumer can provide value to a supply chain by exclusive provision of personal data,  
28 to allow a supply chain to create specially customized and targeted products and services.  
29 However, it is not necessary for a customer to have an exclusive contract with a supply  
30 chain (although this might bring greater loyalty benefits and payments).



1 8) Methods within SDI to recommend information to reveal to a supply chain, to  
2 maximize the value of a loyalty bonus offered and ultimately received by an  
3 agent. Of course, information cannot be falsified, but some information can be  
4 withheld.

5

6 Applications of “bribing” the consumers of a vendor:

7

8 • A commercial entity is sustaining bad publicity from an article published in  
9 newspaper X. The company could offer the customer base of the publication  
10 (identified via SDI) an incentive to temporarily disrupt or boycott the publication  
11 until corrections/changes are made to the article. The company could also offer  
12 discounted products or services to the consumers, based on information (from  
13 SDI) about their consumption patterns. A more indirect discount could be  
14 provided via other vendors further up the supply chain, via multi-vendor  
15 exchange.

16

17 • Company A is a high tech start up selling software products and services to  
18 established vendor B. Vendor B identifies the substantial value proposition of the  
19 products and services provided by the start up, and decides that it should provide  
20 the services for itself “in-house”. Start up company A can go to the consumers of  
21 vendor B (information from SDI) and offer a percentage upside, e.g. percentage  
22 equity in the start up, if they will apply a threat to boycott the vendor unless the  
23 vendor agrees to do business with the startup. Company A might even identify  
24 consumers or suppliers further down the supply chain, and apply indirect pressure  
25 to B.

26

27 • Company A is developing a new commercial initiative which may be competitive  
28 to that of another smaller company B, and might harm B’s market share. Assume  
29 that consumers prefer B’s products to the future products of A, and that B can  
30 make more profits than A. Company B could convince the customers of company  
31 A of these facts, offering them a percentage of potential upside and/or discounted





vendors of its own predicted value, while withholding information about other vendors.

- SDI can seek coalitions of vendors with similar interests to combine their bargaining power, and attempt to compel vendors to enter particular preferred value chain scenarios.

SDI is privy to most or all of the information relating to all of the companies in a supply chain, and can be trusted with the task of creating strategic plans between different companies, potentially making the supply-chain more efficient and benefiting all vendors. The strategic plan might suggest:

- 1) New business relationships between the most synergistic entities
- 2) Contracts for existing commercial relationships which embody and enforce specific terms of those relationships within the framework of this optimal supply chain system.

SDI might also create a detailed system for accounting for the effect of a contract on the rest of the supply chain (individually and collectively), and predict short and long-term effects of contracts in advance. This information can be used to encourage the desired contractual activities on the part of vendors in a chain. There may also be cases in which the direct/indirect economic impact of a contract upon a vendor results from an additional secondary economic advantage to that vendor besides the trickle down economic effect; for example because of strategic commercial benefits. These factors can also be considered in an economic benefit model.

Given models to compute the comparative economic value of proposed contracts to vendors in a supply chain we can demonstrate a “trickle down” economic value chain in which the success of company A is directly related to the success of company B. This information can be presented to company A and each intervening supplier. SDI can then establish long term contracts between each entity in the chain to insure that each

1 company receives the benefits of products and services from its other suppliers which are  
2 of greater value than without the contract. This trickle-down benefit can be used to entice  
3 a company to enter into a recommended contract, for its own benefit and the benefit to  
4 company B.

5  
6 The benefits which each vendor is predicted to receive in a proposed value chain can be  
7 pooled together and used to “bribe” each vendor within the value chain to enter into the  
8 preferred contract. While some value chains may actually be willing to provide greater  
9 compensation to a vendor, the goal is to link a vendor’s decision to the value of the  
10 *complete* value chain, both future and present.

11  
12 Of course, vendors may face risks in committing to long-term contracts, particularly  
13 when many vendors in the chain must prosper for the value to be realized. SDI might  
14 introduce a number of decommitment clauses to allow a vendor to back out of a plan,  
15 perhaps reverting to a “plan B” which protects the interests of harmed vendors while  
16 retaining as much value to other vendors, i.e. allowing all vendors to contribute to the  
17 cost of failed contracts. Alternatively, we might create a *futures market* in which the risk  
18 of the future value can be traded in *real options*, within a public market—providing a  
19 financial instrument to share risk. In this case, if some measurable form of economic  
20 benefit such as sales value (**for e.g., the particular relevant product line**), possibly  
21 stock values, profits, etc. falls below a target level the vendor may be compensated for its  
22 initial risk, i.e., secure a worse-case outcome.

23  
24 It may be important in high-growth markets such as technology, with high risk  
25 companies, to allow insurance for contractual non-performance of suppliers.

26  
27 SDI can leverage available data to compute the most beneficial economic scenario for all  
28 associated vendors collectively, and also to compute optimal scenarios for each  
29 individual entity. SDI can disclose the complete picture to each agent, and then allow  
30 individual vendors to optimize their local plans in a global plan. To effect a new scenario  
31 though, agent agreement is required. Another mechanism could have SDI report the

1 optimal global scenario, and the optimal scenario for each individual agent. Then we  
2 could allow agents to negotiate, to find a *multi-agent compromise outcome* that is more  
3 efficient than the current solution but possibly not as good as the globally optimal new  
4 supply chain. The decision of an individual agent to enter into a proposed contract will be  
5 determined by the cumulative value, short term and/or long term, to that entity, consisting  
6 of the predicted value of taking an action (e.g. breaking a contract, initializing a new  
7 contract) both in terms of immediate bonus to a vendor, long-term individual benefit, and  
8 trickle-down benefit because of global supply chain improvements.

9  
10 Each vendor must provide SDI with as much information as possible about the particular  
11 conditions under which it will perform new contracts, such as its preference between  
12 short term and long term predicted benefits, etc. SDI looks for outcomes that are  
13 favorable to individuals and to the overall system. Value can be realized in short-term  
14 payments and long-term bonuses.

15  
16 *Additional examples include:*

- 17
- 18 • Company A may threaten to file a lawsuit against company B (which could be  
19 very damaging to company B). In a similar barter exchange, company B can  
20 **barter** its own value (in the form of upside or products/services) as offered to the  
21 customers (or vendors of the customers) of company A, targeted to only those  
22 customers that reside within the same supply chain. Thus, by limiting the value  
23 proposition to only these entities, and not to those which reside within the supply  
24 chain of a competitor, the value of the potential upside is worth more because it is  
25 partially predicated upon their own direct success. Again, it is also possible to use  
26 an indirect strategy, bribing entities or individuals of influence within a given  
27 company from which a particular action is desired.
  - 28
  - 29 • In another novel application there are certain activities among certain individuals  
30 which are considered to be mutually and collectively advantageous or  
31 disadvantageous to all members of a certain vendor community. We can have all



contract all the relevant strategic partners before a funding commitment, to provide more information and reduce risk. Other equity shares could flow from marketing agreements, e.g. a co-branded affiliation of all further products/services. The new company might also guarantee an exclusive contract with another entity. An entrepreneur might also be willing to sell a company to the more established entity, e.g., within a year or two, this can be contracted at the outset. The present framework allows value from long-term strategic alliances to be released.

Within SDI we can adapt a user's personalized portal browser to favor the vendors providing such value opportunities to the consumer. A vendor could achieve additional market share by complying with requests of another vendor. The value proposition as presented to the user can leverage the personalized browser user interface to prioritize a vendor's offers in the future.

*Example Commercial Application.*

An investor has invested in a very high risk venture relating to highly speculative stem-cell research. After years and millions of dollars, the venture finally goes out of business (or becomes acquired on the verge of bankruptcy) the investor loses all of her investment, however terms in the original investment contract (which involved at the time the approval of a variety of medical organizations most likely to benefit from the eventual medical science benefits of the research) now enable the investor to acquire in kind benefits (which may include anything from equity in these various medical organizations to goods/services provided by them). A portion of the value may even include a sizable amount of value provided by the original company and/or indirectly vis-a-vie the remaining medical entities, a substantial amount of value within the barter exchange for favors.

With time, stem cell medical applications proliferate thus, the "upside benefit" increases in proportion over time. However, if this "up-side benefit" (of the original collection of

1 contracts from the various relevant companies in the field to which the original research  
2 activities pertained) was insured through the futures market the investor can benefit from  
3 the value and growth. Investors can include researchers whose contribution of time and  
4 effort and dedicated focus on the problem had been estimated. This can allow valuable  
5 talent to receive the incentives to join the initiative from the outset in a way which would  
6 be very difficult to achieve with other standard recruiting approaches. It is important in  
7 this scenario, however not to over incentivize the managers and executives such that they  
8 may lose incentive to make the company succeed. In this scenario that the company does  
9 not go bankrupt, proper incentives could be provided to further encourage similar  
10 scientific and innovative entrepreneurial initiatives in the future, thus a portion (perhaps a  
11 fraction of the upside in these other entities could be provided even in the event of  
12 success (and irrespective of the level of risk associated with the original venture).  
13 Certainly the field of stem cell research is only one example of many speculative new  
14 technology fields with potential for enormous overall impact upon all aspects of the  
15 economy and society. Another example cited elucidates other potential features and  
16 variations of the present novel scheme. The field of na---- technology is believed by  
17 many to become perhaps the most important and pervasive technology paradigm of the  
18 twenty-first century impacting almost every industrial sector. As such, many different  
19 critical problems will need to be solved, each with potentially enormous associated  
20 economic opportunity. There is, however, considerable risk associated with the  
21 probability for industry and research institutions to realize these significant returns on  
22 investment for any given industrial sector or application domain (however, by no means  
23 is this true for the overall field of nanotech). Also breakthroughs in any one application  
24 domain will substantially advance the state of the art for potentially all other application  
25 domains. It is reasonable for SDI to create a futures market in which it predicts certain  
26 long term but very realistic goals for each application domain of nanotech. Through  
27 cooperation with industry and the investment community it may further pre-contract with  
28 various present and potentially future entities which are most likely to be direct  
29 beneficiaries of the technology which is spawned from the initial extremely high risk  
30 commercial ventures pioneering the basic technology. As in the stem cell commercial  
31 example percentages in the various longer term beneficiaries can be used as a strong

1 inducement to present prospective investors in the pioneering commercial entities. As  
2 suggested, this percentage may be inversely related to the degree of success of the initial  
3 venture and if desired, the risk may be further reduced through the use of a futures market  
4 which essentially assures the predicted economic -----marks of these various  
5 prospective commercial endeavors. The wise investor in the futures market will seek  
6 diversification of investment across many different technological sectors of nanotech  
7 (thus assuring the large upside which is certainly to occur universally but not necessarily  
8 for any one given application commercial domain. SDI could further mediate the long-  
9 term opportunities to the original investor by mediating, for example, technology transfer  
10 and licensing intellectual property to present or prospective commercial entities which  
11 have thus agreed to operate under the terms as provided by SDI. In exchange, SDI could  
12 also (as above suggested) provide a resource and technology sharing service and  
13 associated partnership brokering service (between potential multiple companies  
14 containing potential commercial synergies). In addition it may work with the investment  
15 (particularly incubator community) by which it may, without disclosure, of commercial  
16 strategies and technologies which are presently or prospectively planned across the entire  
17 nanotech industry of disclosed suggestions for particular commercial and technological  
18 opportunities which would appropriately complement and benefit the global strategy or  
19 nanotech which is known exclusively by SDI. This approach in addition to emulating  
20 the same united focus and integration of multifaceted technological and R and D  
21 initiatives can also ensure the proper distribution of efforts within each technical  
22 application domain, the appropriate sharing of information whenever potential for  
23 technological complementarity, but not commercial competition exists (which SDI  
24 strives to achieve on an industry wide basis) and the assurance that companies and  
25 departments do not duplicate their efforts if at all possible. The investment community  
26 through SDI could certainly stimulate conditions as above which band aid the adherence  
27 to these terms as a condition to investing. SDI, through incubators could additionally  
28 disclose particular commercial and technologic opportunities which are both synergistic  
29 to the global SDI strategy as well as pre determine (visa vie SDI's extensive in depth  
30 research on the industry to possess significant commercial opportunity. Based upon  
31 some existing companies technical proficiency and commercial strategies, some of these

1 emerging opportunities could be disclosed as potential (suggested extension of existing  
2 departments or divisions based upon their individual areas of expertise and particular skill  
3 sets of its employees.

4  
5 The present system methodology (of business to business and business to consumer) may  
6 also be extended to include other variations including consumer to business or consumer  
7 to consumer. An example of the former could be an employee who is slated for lay-off or  
8 firing within a large organization. The indirect effect of the lay-off could result in a  
9 negative impact upon the individual's financial welfare and indirectly upon the  
10 professional contributions of that employee's children. The likely potential beneficiaries  
11 (corporations) may identify potential lost revenue streams and effectively bribe the  
12 present organization with "value" to avoid the lay off. It is also noteworthy to mention  
13 that the justification for accepting the offer for the second organization is more  
14 compelling if there exists a "trickle down" economic impact, ultimately upon the first  
15 organization (and certainly SDI, within the present framework, would factor such small  
16 economic negative factors into the optimization model for the most efficient supply chain  
17 of vendor relationships).

18  
19 Within the scope of the present example, it is possible to create a form of insurance in  
20 which SDI identifies extremely similar circumstances where similarity is measured  
21 primarily by the nature and scope of the opportunity and its associated risks as well as the  
22 entity whose critical action is necessary to bring about the desired objective. In this case,  
23 the individuals (or entities) benefiting from the desired action are able to effectively  
24 conceive to contribute a certain level of value to the collective group to effectively insure  
25 the desirable outcome on behalf of all of the members of that group, such that, if the  
26 desirable action on behalf of the other entity (in this case the large organization  
27 contemplating substantial layoffs) does not concede to the desired action for one or more  
28 of the individuals, this value contributed from the group is used to apply still greater  
29 economic pressure upon the entity to perform the desired action. If this does not facilitate  
30 the action, the value is instead applied in another (albeit less optimal) form to compensate  
31 at least a portion of negative economic impact upon that individuals/entity of course, the









Thus it is a primary objective within the scope of the present system's formation of common political support strategy groups to ascertain each individual's user's ultimate political objectives based upon the analysis which provides analysis and prediction of the present and likely future effects, both politically and economically from the standpoint of that individual (which may include social, professional and predicted individual investor oriented direct and indirect consequences at a detailed level as it relates to that individual. And secondly, to formulate a global political strategy based upon these objectives which are represented by an SDI agent and which is able to recommend actions and, as such, act as a coalition on behalf of the constituent individuals, and including negotiations for further coalition formation on various portions this global strategy in order to further enhances the collective power wherever common interests can be safely identified and shared to both groups' mutual best interest.

The effects on commercial industry, even specific businesses in which certain actions are likely to result, may be modeled and presented to the user as well as the effects upon the user in light of the investment portfolio of the user. If a particular political position were to be taken, the system could even recommend-----  
-----which is most compatible with a particular political position.

## Continuation to Supply Chain Section

It is also possible to utilize the extensive information in SDI to instruct each political support group's most advantageous supply chain structure in order to be able to recommend certain strategically advantageous cases of tax dollars which facilitate the construction of such supply chains in accordance with the above methods above described. It should be noted that a "political group" may be further subdivided into groups based upon purchase affinities (which themselves suggest different group divisions) which in turn may bribe or boycott commercial entities as needed to achieve their most advantageous strategy as consumers or simply individual users (as economic entities). Each subgroup (via SDI) may further augment their economic leverage with



1 communities (who elect those officials) as well as contribute indirectly via the available  
2 tax base. Or certain commercial entities could, for example, directly or indirectly provide  
3 economic benefit to certain individuals, e.g., if they are employees, investors or  
4 commercial entity which is in some way commercially dependent upon that company,  
5 e.g., for commercial business if it is a customer or alternatively a supplier and the impact  
6 of its resulting success in general (or even local presence) would again provide trickle-  
7 down economic benefits to that entity. In accordance with the present application it may  
8 even be possible to estimate values for such factors as, for example, the economic impact  
9 20 years in the future upon commercial business and industries in the local region (as  
10 well as nationally) if local schools are provided with a 10% vs. a 7 ½% share of the local  
11 tax base. With these predictive models, the associated prospectively affected commercial  
12 entities could, effectively encourage SDI acting on behalf of the local citizens to form a  
13 coalition which makes their voting of the relevant elected official(s) contingent upon the  
14 increased tax allocation for local schools (the specific preferred users could be predefined  
15 as well). SDI acting on behalf of the interests of the prospectively affected commercial  
16 entities could bribe the citizens with either (preferably) such benefits as in kind,  
17 goods/services which could in turn be leveraged through the supply chain or through a  
18 barter exchange system or employment benefits or wage increases (which again could be  
19 effectively treated as a barterable commodity subject to the above conditions of avoiding  
20 interactions with competitive entities. Such modelling could be extended to many other  
21 situations, e.g., allocation of funds to police, fire and EMS services, approving and  
22 providing highway improvement again may provide the above benefits etc. providing  
23 safety and preserving the health, welfare and life (in addition to educational quality) of a  
24 certain number of individuals (each with a certain predictively modeled) local economic  
25 impact. The cost versus impact affecting each citizen could be presented to them by SDI  
26 accordingly in order to elucidate their decision making processes.

27  
28 There may be certain instances in which users at an individual level may be able to provide  
29 predicted -----this benefit can be quantified towards commercial entities or  
30 even individual users if particular events can be achieved (or particular events prevented).  
31 For example, consider a student who is determined to be capable towards achieving



1 Information which is enclosed may, for example, contain in addition to the identity of the  
2 prospective investors, the complete discussion of the economic benefits which each  
3 individual investor could potentially sustain, however, this information is disclosed only to  
4 the SDI representing the collective interests of the prospective investor coalition (a primary  
5 directive of which is to never compromise the data disclosure policies of any individual  
6 investment entity, even if such disclosure to one or more of members of the group may  
7 ostensibly benefit the collective group). It is thus the objective of this collective SDI agent  
8 to make individual data disclosure recommendations for each prospective investor that will  
9 optimize the probability that the basic framework of the collective cooperative group is  
10 successfully created while enabling the individual entities to withheld sufficient strategic  
11 information from the group in order to optimize its own negotiating position with the other  
12 entities. For example, it would be particularly advantageous to disclose credentialed  
13 information regarding the benefits which the investor could provide to the collective group  
14 of investors (based upon the existing investor information disclosed by the associated  
15 companies) and information regarding the existing relationships which that investor has with  
16 existing companies which could, in turn, provide positive benefits to the companies  
17 constituting the global investment strategy of the group and (in the case of negotiating with  
18 each given individual investor prospect) the benefit to companies which that investor has a  
19 relationship with and made available to the present investor prospect. These companies  
20 may also include those which SDI has recommended to that investor as part of the global  
21 investment strategy for that group.

22  
23 user's reactions to various events and stimulate the video programming can be captured,  
24 aggregated and accessed by present viewers such video in real-time or asynchronously for  
25 future viewers of non-live content by user profile (or content profile) or by similar users to  
26 that of the profile of the user. Observing reactions/responses to political or ethnic jokes by  
27 that group which is targeted or other groups may be interesting to users to observe as a  
28 simple example. Reactions/responses to news or political events by different groups or  
29 those most effected thereby may be of interest to users.

### 30 1.7 Off-line Variations





1 If a smart card is used this user profile data may not have to be remotely retrieved but may  
2 be stored on local memory on the card itself along with the client-based pseudonym proxy  
3 server. In one novel variation, a card is done away with completely by virtue of  
4 revolutionary technological breakthroughs in being able to instantly and positively identify  
5 buyers biometrically using iris scanning techniques (which may in a variation be further  
6 combined with facial recognition techniques). Many vendors will wish to utilize user profile  
7 data in order to deliver promotions targeted discounts and promotions (see pending patent  
8 "System for Customized Prices and Promotions")

### 9 **1.7.1 Location Enhanced SDI System**

10  
11 The co-pending application entitled "Location Enhanced Information Architecture" (LEIA)  
12 describes an integrated advertising delivery platform which selectively targets buyer  
13 personalized advertising based upon both the buyer's personal profile and the present  
14 location of the buyer which may suggest appropriate ads from vendors which are local to the  
15 buyer, wherein user identifiers (UID's) which could include any of the above identification  
16 media provide the essential elements for this buyer targeting platform. With LEIA  
17 information providers can utilize location information, in addition to static and dynamic  
18 profiling information. The method customizes the information that is displayed on a private  
19 or public information device to the real audience in the vicinity of the device, instead of a  
20 predicted audience. LEIA collects an extremely detailed and comprehensive information set  
21 about the daily activities of a user, enabling enhancement of the user profile with location  
22 information and temporal activity patterns. The co-pending LEIA patent suggests  
23 appropriate application environments, for example in a smart home, an office, on a mobile  
24 shopping device, and in an automobile. A LEIA-based system stores personal information  
25 on users.

26 We can extend LEIA by incorporation with the Secure Data Interchange system that we  
27 teach in this patent. SDI enables the user to receive the benefits of powerful and well-  
28 directed information, but within a system that respects his/her privacy requirements. The  
29 interchange acts as a secure data warehouse for users and information providers, enabling  
30 information providers to target users without revealing private information to the providers  
31 directly. As described in LEIA, the privacy architectures provided for in LEIA and SDI are

all the more critical as components of an SDI enhanced LEIA system architecture given the extreme sensitivity of personal location data which must be securely protected in both real-time and within an asynchronous context.

LEIA customizes information that is displayed to an information recipient based on object profiles and physical location of users. Presents the information most relevant to the REAL audience, not a predicted audience per se. Because of LEIA's ability to combine user's preference information with information about their current physical locations and deliver advertising which is both of personal and location contextual relevance to the user, LEIA provides a valuable enhancement to the targeted advertising services provided by iamworthit. For example, instead of advertisers being limited to accessing desired preference attributes of users in the pseudonymous user database exclusively they may additionally access prospective target recipients by present (or anticipated future) location parameters of users connected to mobile and even terrestrial communications networks.

In accordance with the iambworthit targeted advertising platform either advertisers select user profile attributes which they are interested in or the system matches their ads automatically to the most relevant users. The former application may be performed either within the context of the pseudonymous user database which the vendor queries or targeting rules which s/he provides which control the advertising and pages which are dynamically generated for each user based upon his/her user profile. Whether the user's location is detected on a mobile network or a terrestrial network including telephone or cable TV, LEIA's location features are able to provide additional useful features regarding users to the vendor's rules interface or to the pseudonymous user database controlling dynamic page generation at his/her site.

For example, at a bookstore, we can recommend isles and particular books; at a supermarket, can play music preferences; smart-radio, play appropriate channels in a cab based on target object profiles (as meta-data). As suggested in issued patent “System for Broadcast of and access to Video and Other Data Using Customer Profiles” appearance of relevant selections can be continuously scanned for, dynamically selected and presented to the buyer in the form of “Virtual radio station”. Such a system can also be linked to a service for making an instant purchase, or linked to a database (in conjunction with LEIA) to



Other extensions of this scheme are also considered e.g. within the context of the user's office, or automobile and pedestrian activities. This application may thus extended the usefulness of the iamworthit model to advertisers in being able to target users through the presently anticipated on-line media as well as networked appliances and in either case, based upon the relevant context of users' present activities and behavior (and from this potentially their inferred moods or mental states) within their homes and elsewhere. Additionally (perhaps most importantly), LEIA provides a highly beneficial value to vendors whose customers purchase primarily from the vendor's bricks and mortar store presence. SDI can identify multiple vendors which share a common physical location. Additionally, LEIA can also identify the physical location of buyers on mobile and terrestrial networks and suggest the sharing of off-line prospects for these off-line retailers on the basis of not only interest but physical location.

### 1.7.2 Digital Set Top Box Methods

Similarly, we can deliver targeted advertising and other information through cable TV systems, as described in the issued parent patent application entitled “System and Method for Scheduling Broadcast of and Access to Video Programs and Other Data Using Customer Profiles” US Patent # 5,758,257, and co-pending application entitled “Broadcast & System for reduced memory terminals broadly address the use of cable systems as an interactive medium (in a bi-directional network architecture) for purposes of delivering targeted advertising targeted advertising and other information to the consumer based on user profiles”.

In this system customer behavioral data is collected at the digital set top and the upstream channel enables these profiles to be processed at the lead end server. These detailed profiles may then be subsequently transmitted down and stored at the level of the individual set top. The cable environment is a two way interactive medium. The bandwidth allocation is inherently asymmetric. Separate channels can push parallel adverts, which are selected at the set-top-box according to a buyer's profile. Each channel can have associated meta-data to allow matching at the set-top-box. As an alternative variation, full motion advertisements may be down loaded in the form of applets to the digital set top box and displayed to the

1 buyer in similar fashion as described above. Or full-size commercials or infomercials could  
2 even be downloaded and inserted in place of existing commercials on the video stream.  
3 Digital tags used for queuing for ad insertion technology enable this capability. The  
4 preferred commercial implementation of the system is used within the context iamworthit  
5 (the buyer infomediary service in which the viewer receives value for advertising received).  
6 This method is preferred (in as much as in the following alternative, the cable operator does  
7 not need to be in the loop). With that said, in an alternative embodiment, however, requiring  
8 the cooperation of the cable operator various commercials could be broadcast in parallel  
9 during commercial breaks. The user profile at the set-top is deployed to select the most  
10 appropriate commercial accordingly.

11 User profiles regarding viewing behavior can be collected with relatively minor up-stream  
12 transmission, e.g., to infer whether a buyer is viewing a program the set-top may transmit  
13 the viewers current viewing selection, e.g., two minutes after the beginning of each half hour  
14 and two minutes before the end of the half hour (and possibly at additional intervals during  
15 the viewing segment as well). In the point-to-point access and delivery of personal  
16 information as well as the passive monitoring of viewing behavior (for profile generation at  
17 the head-end and subsequent user profile delivery for the set-top) because of strict buyer  
18 privacy regulations in the cable/satellite industries the use of the pseudonym proxy server  
19 architecture (as described in the parent case and integrated into the above referenced case) is  
20 extremely important.

21 Once interactive television is enabled by access to very large amounts of bandwidth, it will  
22 be very feasible from a bandwidth utilization standpoint to transmit more interactive  
23 content. Predictive caching is still extremely advantageous in as much as advertising (unlike  
24 interactive video real-time on demand video content) tends to be controlled by targeting  
25 rules of the advertiser rather than the user who receives value in exchange. Of course, radio  
26 or music content may also be delivered in conjunction with preloaded audio commercials.  
27 And any of the above content is deliverable over the same networks to other appliances such  
28 as the P.C.

### 29 **1.7.3 Optical-Based System Extension**

30 Iamworthit in a future version of the system could incorporate a novel optically based  
31 medium for delivery of the advertising. The protocol would enable the delivery of user

targeted advertising in a non-electronic environment directly to the user from a static source in visible contact of the user. The system incorporates the following aspects: The user wears glasses or contact lenses which contain a Polaroid film in which the optical medium is polarized for only certain very narrow and specific wavelengths. Advertisements (or other messages) may be presented from signboards which display from the same surface ads appropriate for every different user profile segment for which a unique message is appropriate. The Polaroid film for each user is unique to the user segment to which they belong. Each user segment's optical film filters out all of the particular wavelengths of all other messages except that which is appropriate to the user segment to which that user belongs. Alternatively, the current application may also be relevant to virtual tags as users may provide instead optical messages specifically targeted by a user to other users characterized by their profile features which establishes the entire criteria for the user segment they fall into for which an appropriate message may be targeted.

## 2. Interactive Data Mining Applications

In this section we describe interactive data mining applications, which involve dynamic two- or multi-way communication between agents. For example, within the context of a vendor-consumer interaction, the key difference is that information flow is not one-directional. An important example is time-of-purchase, in which the consumer's agent initiates a request for a response (with a counter bid) from a number of vendors. The vendors can access profile information about that consumer, and then make personalized responses. In general, this bid-response processes can iterate, with vendors/consumers exchanging information over multiple rounds. The exchange of information may occur within SDI in initial stages, with the system acting as a trusted intermediary until the agents are ready for an introduction.

We describe time-of-purchase competition, and then move on to general match-making applications, which allow interested parties to exchange information by mutual consent (as indicated in SDI rules), with initial information exchange autonomous. There are a number of interesting match-making applications, for example within a system for sensitive negotiation, introductions. The system of SDI can play the role of a trusted intermediary, so that only suitable introductions are

made. A vendor can place a request for a certain type of meeting without its competitors knowing that it has made such a request. Finally, we describe a “real-time experts market”, in which experts can respond for payment to questions placed by other agents. The system of SDI allows a useful matching of agents and experts, and also allows data mining to check on the feedback a particular expert has received from previous clients. We also suggest an application of interactive SDI-based data-mining to resale markets, and to a decentralized transportation management system.

## 2.1 Time of Purchase Competition

One application of the system of Secure Data Interchange is in “time of purchase” competition, where by a networked system of vendors can subscribe and receive the opportunity to place counteroffers to users that are about to make a purchase. This application is enabled within SDI because the user agent can remain anonymous while receiving counteroffers, but still use the system of SDI to provide profile information that can allow vendors to make attractive personalized offers. The client-side SDI proxy notifies a central SDI “time of purchase” server, which can:

- Determine appropriate vendors to prompt with an opportunity to make a counteroffer.
- Create a one-time anonymous identifier for the user, linked to an appropriate set of profile information, and allow vendors to execute queries under that profile.
- Collect counteroffers, and pass them to the user’s client machine.

We also describe how to generate coalitions dynamically, based on the privileged position of the time-of-purchase server in the marketplace. The time-of-purchase server can transparently bundle purchases from disparate end-users, and negotiate quantity discounts with vendors. Buyer Coalition formation can significantly improve negotiating leverage as the collective buying power of numerous individuals using the time-of-purchase competition feature, and be made completely transparent to users.

Time-of-purchase allows a user to elicit dynamic market competition between vendors. For example, when a user browses a book at Amazon.com if she/he chooses to activate the time-of-purchase competition feature from the browser, a notification will be sent-out in real-time to all competing book stores which also offer that product section. The notification contains





1       b) The vendor is un-cooperative, and the client-side SDI proxy needs to use other  
2       techniques to classify the type of product or service that a user is requesting from  
3       a vendor.  
4

5       The central SDI data warehouse maintains a searchable index of vendors that provide  
6       certain products and services, and sending buyer-profile and purchase information to  
7       appropriate vendors just before a purchase is made. Vendors can register with SDI to receive  
8       time-of-purchase competition opportunities, and can also register to allow time-of-purchase  
9       competition with purchases on their web pages (case a, above). Of course, it is unlikely that  
10      a vendor would do this without incentive, perhaps one incentive could be a discounted price  
11      for registering to receive opportunities for time-of-purchase competition.

12      In fact, this problem of client-side identification of transactions is a key problem that must  
13      be addressed in a description of client-side user profiling and the submission of information  
14      to SDI. The system of SDI requires a method to know what the user is doing. Click streams,  
15      typed query information, provided profile information etc.

16      A vendor that registers to receive time-of-purchase competition opportunities agrees on a  
17      standard classification system for products and services, and to state what types of products  
18      and services it wishes to enable for time-of-purchase competition. The classification system  
19      encodes the product or service that a user requests, at an appropriate level of detail to allow  
20      other vendors to make reasonable counteroffers.

21      The central server maintains a database of vendor interest sets, so that it can  
22      determine from a classification code which vendors will be interested to provide  
23      counteroffers.

24      For example, if the product is "flights to the UK" then the code might encode the dates that  
25      the user wants to travel, and his/her preferred departure and arrival airports.

26      Take case (a) , where we assume that a vendor is cooperative and provides the SDI proxy  
27      agent that implements time-of-purchase for the user with a final product/service code and  
28      price. The system works as follows:

- 29       1. User enters into a dialogue with vendor.
- 30       2. Vendor makes "final offer" to a user, and the user is about to make a purchase.

31      The vendor has also informed the user of the product/service classification code

- that encodes relevant information about the type of product the user wants to purchase.
3. User clicks on ‘time of purchase competition option’ on its local SDI-enabled client.
  4. Client sends message to the SDI time-of-purchase competition server with (User profile information/identifier, one-time time-of-purchase ID, product/service classification code)
  5. Time-of-purchase competition server looks up the product code, and sends the user profile identifiers and a time-of-purchase ID to relevant vendors.
  6. Vendors can construct competitive offers based on profile information accessed in the central SDI data warehouse, and information about the existing offer.
  7. Vendors send the counteroffers to the time-of-purchase server, which forwards the offers to the user’s client-side SDI proxy.
  8. The user is presented with a set of alternatives, and can make a final purchase decision.
  9. The final purchase is completed with the appropriate vendor.
- We allow other vendors to access profile information about a user based on the profile information that the user provides to time-of-purchase, i.e. based on the profile which it wishes to use for the purposes of having appropriate counteroffers constructed. At this stage the client-side SDI proxy agent might also send additional profile information to be stored within the SDI data warehouse.
- It is useful for vendors to access profile information about a user (anonymously or pseudonymously, as defined by the user’s proxy agent) because products and services can be configurable goods or services, and a vendor can compete on more than price alone, but also in other dimensions.
- The one-time time-of-purchase ID is constructed client-side in step (4) to act as a temporary identifier for the purposes of soliciting competitive responses from vendors. Vendors can be prevented from making direct offers to users, all counteroffers must be forwarded through the time-of-purchase server.



4). **B2B Commerce** – Numerous commercial applications including standard Web-based environments for B2B e-commerce including buyer-side and seller-side e-commerce applications, vertical B2B portals and trading floors.

5). Job Positions – Employers posting job positions may be notified by the service of competing job offers which employment candidates are browsing ( and ultimately accept if this information is indicated by the candidate). Resumes and, if desired, other profile information is provided to all competitors. Depending upon the profile of the candidate and/or competitive employment opportunities which a particular candidate is observing (or has recently observed), employers may wish to customize offers dynamically on an individual basis.

#### 2.1.4 Integration into a Shopbot Interface

We also suggest the integration of time-of-purchase competition into a standard shop-bot interface e.g. with comparative features of the products associated with the various offers across a variety of product criteria, or at the least price can be used to provide the user within simple recommendations. With this, even without a time-of-purchase offer the user can at least compare its offer with fixed-price deals in the wider market place. We might also provide a user with historical information about previous (low) price offers, within a Yahoo-like portal for e-commerce; to provide a user within information about a reasonable price for each individual purchasable; i.e. the very lowest price that each given item had been offered to a customer previously using time of purchase competition. This site could also act as a stand alone proxy server (like anonymizer) which inserts this information by overlaying it on Web pages as the user browses the web.

The time of purchase competition feature is designed to be accessible to a user within the context of or most any information access mode relating to browsing on the Internet e.g. Web or portal browsing, receiving email or “push” content, submitting queries (e.g. specific names, categories and/or desired features (or combinations of the above) of products which the user identified as being of present interest.



1 Time-of-purchase enables vendors with competing products or services to receive automatic  
2 notification when a buyer is about to purchase a relevant product or service. A vendor can  
3 also receive information on the profile of a buyer, and the offers made by other vendors; and  
4 submit counter-offers to a buyer via the buyer's SDI-enabled client. The buyer can then be  
5 presented with a final set of offers, before making a purchase decision.

6 A buyer can also configure its profile management rules within time-of-purchase to provide  
7 profile information relating to the buyer's sensitivity to discount offers, customer loyalty  
8 with other vendors, value responsiveness (bargain driven), responsiveness to high quantity  
9 discounts (for only those categories which the buyer makes frequent or large purchases), etc.  
10 This information can be *certified* by the buyer's client-side SDI proxy, as a fair  
11 representation of the purchasing habits of a buyer, and can be aggregated across product  
12 types to protect a buyer's privacy.

13 The SDI time-of-purchase server can identify vendors with similar products or services,  
14 either using a static index which maintains vendors in particular product domains, or  
15 through dynamic profile matches between the target object profile of the web site that the  
16 buyer is currently browsing and target object profiles of the web sites of other SDI-enabled  
17 vendors. Alternatively, perhaps classification and clustering techniques can be deployed to  
18 identify similarity between vendors at the level of target objects, i.e. the products that a  
19 vendor offers; as indicated by virtual tags provided by a vendor that profile its products and  
20 services.

21 The time-of-purchase competition system may also draw in static "listed" offers on a  
22 vendor's web site, to prevent a vendor over-charging a user. In addition, SDI might track  
23 purchases and via clustering techniques request prices on similar or competing products to  
24 the product that a user is about to purchase. These products and/or services may also be  
25 released to a user. When presenting a choice of products and/or purchases to a user we can  
26 allow many features to be summarized, for example features which relate to the location of a  
27 vendor, terms of shipment and insurance, and user ratings.

28 Vendors are notified, and provided with the ability to access the profile of the buyer (as the  
29 buyer deems appropriate), either with client-level processing or through the release of an  
30 anonymous profile to the vendor. Vendors typically will wish to construct offers through a  
31 rule-based engine, data-mining techniques, or automatic collaborative filtering techniques,

as disclosed in co-pending patent application "System for Automatic Determination of Customized Prices and Promotions" and U.S. Patent #5,754,939, "System for Generation of User profiles for a System for Customized Electronic Identification of Desirable Objects" as such techniques may be deployed by the vendor directly or via the Secure Data Interchange representing the interests of the vendors.

*Choosing an Offer to Make*

User profile information may include a temporal profile of the buyer's present activities, including search terms, recent page navigations, what pages is the buyer observing presently (and the profile of this page) or even his/her present physical location as well as the general user profile. Any portion of the above particularly the latter two may of course be withheld from the vendor). Or the buyer may be more generous in providing this information about the various static or dynamic features.

Vendors can target buyers on the basis of their preferences and interests, and also within the temporal context of when they are most likely to be receptive to offers (thus providing also a benefit/service to the buyer as well as enabling the competing vendors to increase the price point at which their competition begins, knowing that the buyer is in an optimally receptive condition to accept offers for that item). For example, a user might release to a vendor the current context of a user's purchase, such as the click stream over the past 5 minutes.

In the preferred implementation, vendors are also provided with a (client or web-based) rules interface which enables the vendors to input pre-stated rules with which the system may solicit and respond to competitive offers automatically. In this way data mining may be performed in order for the vendor to determine what a user (or all users) by attribute, tend to best respond to by product, product feature, features or services of the vendor, price, etc. If pre-stated rules are used to automatically respond to a notification with a competitive offer, the nature and degree of discount is typically determined in accordance with the nature and degree of the original or previous offer and/or the user profile as disclosed by the client-level proxy/server to that vendor. In lieu of manually entered rules, co-pending patent application entitled "System for the Automatic Determination of Customized Prices and Promotions" another similar algorithmic methodology may be used as an aid by the vendor in order to automatically determine a competitive offer (or subsequent responses thereto) as well as an aid to the vendor in selecting optimal rules.





1 by the user profile as well as the offer response pairs up to that point in the negotiation  
2 process.

3

4 It is likely that vendors will not compete on price alone, but rather through added-value  
5 services such as offering loyalty bonuses, cross-sells, and two-for-one offer and added  
6 features as well as service advantages such as fast delivery, guaranteed service/maintenance,  
7 warranties etc. Vendors will choose this mode of selling to prevent simple price-  
8 comparison at the client. Conversely, in some instances vendors may attempt to eliminate  
9 the features in order to create the perception of a better deal through marginal price  
10 reductions, then possibly add these features as counter offers to other vendors in which the  
11 same price may be maintained.

12 Accordingly the above referenced algorithm as provided by the customized prices and  
13 promotions scheme, may offer valuable competitive insights to the vendor in being able to  
14 assemble more appealing cross-sells/up-sells and more complex offers (including multiple  
15 items which are adjusted by the system to be extremely desirable to the buyer) more  
16 effectively than the competition. Therefore the client will receive offers from multiple  
17 vendors, and after initial filtering of the large collection of offers, present a choice set to the  
18 buyer.

19 SDI can also present a buyer with ratings or annotations for items across multiple competing  
20 products, for example, within the context of buyers whose attributes (which may be  
21 important pricing features) are similar to that of the buyer. SDI might also monitor previous  
22 offers made by vendors, so that if a better offer was made in the past a buyer can request the  
23 lower price, and try to bargain with the vendor. Typically the resulting purchase  
24 recommendations are presented as suggestions on the user interface to help the user's  
25 assessment of valuing the offer as a "good deal" as well as to accurately represent his/her  
26 own preferences. Some users may, however, use this feature as part of a fully automated  
27 (agent mediated) buying process which is performed in asynchronous fashion (as described  
28 below, this typically is only necessary if the time-of-purchase system is deployed in order to  
29 create buyer coalitions prior to negotiating with the vendors. Of course the time of purchase  
30 competition enabled buyer will distribute the offers in order to elicit vendor competition.



1 iamworthit infomediary service may apply pressure upon the vendor such as creating a  
2 buyer coalition and/or making a believable threat (as described below) which has the ability  
3 to deny future business to that vendor. If the vendor is able to present an exceptionally  
4 attractive offer the buyer may instead wish to stipulate that the offer is valid only if the time  
5 of purchase competition feature is not used by the user. iamworthit's statistics based price  
6 prediction scheme can then predict if this decision would serve the user's best interests or  
7 not based upon the user's purchasing objectives.

8 Furthermore, we can allow vendors to offer payment to a client in return for displaying an  
9 offer to the buyer, and vendors can also bid for space on the buyer's web portal which is  
10 often represented as a profile associated with a pseudonym in conjunction with a description  
11 of the ad space. The purchasing decisions of the buyer may be performed by an electronic  
12 representative of the buyer's wishes (as "buyer agent") implementing the techniques of  
13 pricing/promotion selection algorithms completely autonomously on behalf of the buyer.  
14 However, the best offer can only be presented to a buyer to the extent that the SDI client  
15 level software understands a buyer's model of "value", and can make appropriate tradeoffs  
16 between product features and price (as implicitly inferred by the system through the above  
17 suggested techniques or explicitly stated by the buyer in advance). Nonetheless, this is a  
18 hard problem, and we expect that the buyer will often need to make a final product choice  
19 decision which could be a default in which the system would defer judgment to the buyer if  
20 it's statistical confidence as to the buyer wishes below a certain threshold.

21 The collaborative filtering techniques described in pending patent "System for Automatic  
22 Determination of Customized Prices and Promotions", can allow a buyer's client-level  
23 proxy server, termed the buyer agent in this section, to automatically analyze offers. The  
24 system can also be used to send initial offers to vendors, on the basis of historical  
25 information about the transactions that have been performed between other buyers (which  
26 include the benefit of his/her complete user profile data) and the vendor. Offers can (of  
27 course) be sent to a vendor and its competitors. Finally, after offers that are received from  
28 vendors are pre-screened, they can be automatically ranked for value—using a combined  
29 quality and price metric (again judged within a collaborative filtering framework). Buyer  
30 feedback e.g. average ratings and annotations from previous buyers may form an additional  
31 criteria. The goal is to leverage the database of other offers that have been accepted by

buyers in the past, and form a model of vendors, to determine whether or not a buyer has received good offers (i.e. we can exchange information within the system of Secure Data Interchange, and making more information available increases the efficiency of the market). Offers can be filtered and presented to a buyer in rank order.

### **2.1.5 Time-of-Purchase with Non-cooperative vendors**

We suggest a simple technique to allow a client-side SDI proxy to automatically detect that a purchase is about to be authorized, and the details of the purchase. The model is that SDI, or a third-party SDI client, provides a banking service for a user. The advantage of such a service is that in making a purchase, we can require that the vendor provides information to the bank about the purchase that is about to be made, as a condition for validating payment.

The system works as follows:

1. User interacts with vendor, and decides to make a purchase.
2. User provides payment method, for example an SDI bank account number.
3. The SDI bank account server demands information about the product which is about to be purchased before authorizing payment for the service.
4. The vendor provides the SDI bank with information about the purchase.
5. The SDI time-of-purchase system can now step into the transaction, and solicit competitive offers from other vendors.
6. The user is presented with the opportunity to select an alternative offer.
7. The SDI proxy-agent reports the user's final purchase decision to the appropriate vendor, and denies purchase to the initial vendor if the user accepts an alternative offer.
8. The final transaction is executed.

### **2.1.6 Automatic Buyer Coalition Generation**

The time of purchase competition system can also be used to allow the automatic formation of user coalitions. Coalitions can be generated dynamically based on the privileged position of the time-of-purchase server in the marketplace. The time-of-purchase server can transparently bundle purchases from disparate end-users, and negotiate quantity discounts



The time-of-competition server can also take a position in the market, identifying bundled purchases over time, essentially aggregating demand and making offers of a guaranteed stream of purchases over a period of time based on historical information, again for a discount that can be seamlessly passed onto future purchasers. If the server is successful in taking a position in the market for some good that its user base seem to be interested in purchasing, then in addition to returning offers from vendors in the system, the time-of-purchase system can provide good prices to its clients.

Buyer coalitions typically are able to exert significantly greater pressure upon vendors the larger they are (this fact applies as well to multiple company b to b buyer coalitions). As such iamworthit may also notify other members which in accordance with their user profiles are likely to have an interest in a particular item (or items) in which a coalition of members have demonstrated interest in order to increase the size of the coalition. Similarly, the aggregate user profile of all members of the coalition may be used in conjunction with collaborative filtering in order to recommend those items which collectively best match the preferences of the coalition collectively and thus are able to provide the best deal for the most items possible to that particular coalition. Thus the coalition model is extended to placing competitive pressure upon vendors across multiple items. iamworthit must also

1 utilize the above modeling techniques in order to establish pricing models based on the  
2 value of these different items in light of the total monetary value which the coalition  
3 represents to the vendor collectively. In addition to inferring buyer interest entirely  
4 passively, it is possible for buyers to manually edit their profile, e.g. by modifying or  
5 actively selecting categories of purchasables of interest which she/he would like to be  
6 notified of (e.g. via pager) if/when a buyer coalition for a purchasable of that type is being  
7 established. It is further possible to increase the size of buyer coalitions by allowing a  
8 certain amount of time for buyer's to respond to offers to join that particular coalition (e.g.  
9 24 hours). Typically a couple of days is all the time that is required to achieve a  
10 substantially maximum buyer response. Also, because some vendors may wish to be  
11 (automatically) informed if/when substantially large coalitions approach the vendor, there  
12 may be an additional marginal price advantage by providing an additional period for  
13 vendors to reply.

14 The buyer may wish to allow the iamworthit time of purchase vendor competition system to  
15 operate in fully automatic mode whereby certain "acceptance parameters" are provided to  
16 enable the buyer agent to act fully or autonomously on behalf of the buyer's stated (or  
17 approved) parameters. This technique is likely to provide a central function particularly in  
18 most business-to-business iamworthit commercial applications. The business-to-business  
19 commercial domain may exemplify a couple of other novel features (which are not  
20 exclusively limited to business-to-business). For example, because replenishment of  
21 products or supplies suggest much greater repetition and thus predictability, it is sometimes  
22 useful to buyers to also leverage as part of the negotiating process an advanced commitment  
23 of buyer loyalty over an extended period of time. In the unlikely event of the buyer  
24 changing his/her vendor loyalty after considerable value has been provided to the buyer, the  
25 vendor may control the right to acquire certain assets of the buyer (e.g. corporate stock) as  
26 pre-agreed collateral for the vendor. Also, in the business-to-business domain iamworthit  
27 acts essentially as (or in conjunction with) a buyer-side portal to a web-wide vendor  
28 resource, though it could also be integrated as part of a seller side-portal interface tool as  
29 well.

30 As an additional service to vendors SDI can provide enhanced profile information,  
31 aggregated from other vendors, to enable vendors to provide better focused offers than can



1 be provided on the basis of the profile information directly associated with the pseudonym  
2 of a buyer. Certain portions of the user profile data that is unavailable for direct collection  
3 by the vendor (such as information that is collected on other sites including, in particular,  
4 competitive vendor sites) may reveal important information which enables the vendor to  
5 better target that buyer. As such the secure data interchange representing the collective  
6 buyers may aggregate, analyze and sell this data to the vendor so long as the release of such  
7 information does not negatively affect the predicted pricing or value levels for that buyer as  
8 performed by the above type of pricing algorithm.

9 An interesting variation of this example involves the situation in which multiple entities  
10 represented by SDI users may be negotiating with the same other entity (or individual). It  
11 may be for the same purchasable. Alternatively, it may involve different purchasables or the  
12 represented group may even be a combination of users and sellers who happen to be  
13 presently dealing with the same entity. Accordingly it is important for SDI to achieve an  
14 understanding as to the individual negotiating parameters as well as an assessment of a  
15 market demand model which characterizes the needs and objectives of that entity with  
16 regards to the particular prospective transactions being negotiated which SDI is mediating  
17 on behalf of each party which is captured through the main SDI server. As such it is often  
18 possible to thus attempt to predict the minimal acceptable terms of an offer which would  
19 meet those conditions. The negotiating leverage from all parties is based upon terms which  
20 affect multiple (or perhaps all) of the deals collectively. SDI thus has the power to  
21 significantly strengthen the negotiating leverage based upon a market demand model which  
22 characterizes that entity thus benefiting the entrusted parties on the whole. The statistical  
23 data used for market demand models from the user-centric or vendor-centric SDI is  
24 collected about the other parties which the parties are negotiating with. One useful source  
25 for estimating this predicted market demand curve which they represent. For example it is  
26 possible for end-users as a coalition of users to be automatically formed for this purpose.

27 A user-centric SDI is then assigned to the coalition in order to (a) ascertain from user profile  
28 information (in addition to subsequent active query or recommended offer feed  
29 back/approval information about terms which the coalition would likely wish to receive  
30 from the vendor (subject to negotiation) (b) ascertain from vendor behavior relating to past

users and coalitions terms which would be reasonable to expect from the vendor in order to anticipate reasonable offers/counter offers to present to the vendor.

SDI's predictive tool suite can be used to predict the most likely prospects for a given coalition, i.e., matching users who have explicitly indicated an interest or have engaged in negotiations or past dealings with a given vendor. As indicated the recommended users may be targeted either with the same items or completely different items, the latter case in which SDI will try to negotiate a lower price based upon the aggregate sales from the coalition, not the volume of sales for one particular item per se, as in the first instance. In the present application, however, because negotiating pressure is being applied by a collection of entities upon (typically) one entity (or possibly alliance of entities), an SDI (which may typically be spawned from the main SDI service) representing the coalition is created in ad-hoc fashion. The other entity may accordingly wish to acquire its own SDI to also represent its commercial interests. Each SDI may represent multiple parties. No SDI must of course ever represent an entity which is negotiating with another entity which it represents.





these complex vendor synergies exist even indirectly via one (or multiple) removed vendor-vendor and user-vendor relationships.

#### 2.1.8.2 Coalitions Created For Purposes Of Eliciting Change On The Part Of Vendors

In addition to applying market pressures, one other example in which entities may be automatically introduced into coalitions using SDI may involve SDI automatically identifying certain situations or complaints from users or entities in which there may be common ground for these parties to collectively apply political or legal pressure upon another entity from which the problem arises, or in one example to apply legal tactics for purposes of reclaiming damages from the infringing party. For example, members of SDI could be asked to disclose particular instances of these types of legal infractions even if they are relatively minor (e.g. simply a short fall in acceptable quality of misrepresentations through misleading advertising etc. may be relatively minor inconveniences and/or a site may be provided as well for non-SDI members (e.g. [www.class.action.com](http://www.class.action.com)).

A similar application tailored specifically to the Federal Government could also be provided (e.g., [www.governmentindustries.com](http://www.governmentindustries.com). In one example, SDI's ability to provide annotations excluding complaints and ratings pertaining to particular organizations and individuals is complemented by the current scheme by introducing the ability of users to apply pressure to the organization or individual as a coalition in order to more effectively motivate the desired change accordingly. To this end other tactics may be used, for example, a threat to collectively sponsor negative advertising or editorials. In fact, the above technique for presenting a believable threat could be utilized in this case and backed up by an insurance policy which provides in addition to the guaranties that the threatened action will be carried out, for example, by X date unless a particular condition is met, also provides sufficient funds to follow-through accordingly and in one variation further instill the motivation on behalf of the parties by actually insuring the outcome of the litigation (where the threatened (and guaranteed) consequence is sufficiently great that the entity is much rather proved to settle thus reducing exposure significantly).

Within the present system SDI provides the framework by which appropriate users and sellers may be matched together. It also enables a methodology by which the user interests

are protected through the use of matching of sellers offers to competitive vendors (using  
iamworthit).

### **2.1.7 Single-user Time-aggregated Purchases**

Another extension is to propose a special interest-bearing bank account which we set-up for  
the user which is tied into a wallet or debit/credit card for off-line transactions. We can  
allow a user agent to bundle its own purchases over time, and use the system of SDI to  
guarantee future payment to on-line vendors in return to a good current offer. This is novel,  
because it is like participating in a “discount scheme” without purchasing the right to future  
discounts up front. A user can still invest money that is earmarked for future purchases with  
a vendor.

We describe a wallet application for this purpose, that allows a user of SDI to commit  
currency to vendors for future purchases. The time-of-purchase competition system can  
allow vendors to see the pseudonymous purchasing profiles of users (e.g., over the past 12  
or 18 months) which are relevant to their respective commercial venues and competitively  
bid for the opportunity to acquire from that user a “pre-commitment” for similar spending  
amounts within the same time period going forward (provided that that vendor offers the  
same/similar categories from which the user had purchased that amount in the past from  
another vendor(s).

The “committed currency” in the user’s account is not actually spent and thus remains  
interest-bearing for the user, and because the dollars are exclusively and irreversibly  
earmarked, the actual debit could even occur substantially after the actual purchases  
ultimately occurs with that vendor. Vendors marketing the service may be able to exclude  
competitive offers (as is the case with the standard time-of-purchase competition  
application).

### **2.1.8 Special Vendor Treatment in a Community Dollar Scheme**



1 can be delivered to the user via either the advert replacement system or a pop-up  
2 window/java-script window.

3  
4 Vendors might even compete to "buy" the rights to exclusively deliver competitive offers  
5 to any offer which the user may happen to view and which is competitive to another  
6 product/service which the vendor provides. The terms of the rights which the vendor  
7 actually purchases may include the number of competitive offers which the vendor can  
8 provide in the future, the degree of the conditional "value" which the vendor would  
9 provide through the competing offer, the category(s)/domain(s) which the vendor  
10 purchases the right to compete, the degree of exclusivity/non-exclusivity. A vendor may  
11 also wish to constrain the ability of potential competitors to submit counter bids.

12  
13 Vendors in all other product categories can retain the ability to compete for any offer which  
14 the user receives. It is possible that vendors with exclusive rights may also try to lock a user  
15 into *future* purchases. Again, this ability to lock users into future purchase commitments  
16 may be auctioned to competitive vendors. The dedicated currency auction web site (e.g.,  
17 [www.creditauction.com](http://www.creditauction.com)) is architecturally very similar to the time-of-purchase competition  
18 variation of the service with the exception that it is a market place in which many users may  
19 submit their profiles with request for offers from vendors across any or all categories  
20 relevant to their profile as requested or agreed to by the user. A variation of this dedicated  
21 currency involves a scheme for time shifting into the future the transfer of funds from the  
22 users account to the vendor's (thus allowing the users money to accrue interest during that  
23 period). Thus dedicated currency whether it proceeds or preceeds actual purchase events  
24 guarantees the vendor sales which can be used -----plan capital improvements, attract  
25 credit or investment funding while providing benefits such as -----  
26 and/or interest on money spent to consumers.

27  
28 This reverse auctions for time-shifted purchases may also encourage vendors to form  
29 virtual retailer communities, to accept a common currency and offer a full range of  
30 products and services to a user. Any offers involving the user dedicating any form of  
31 "value" for the user for future purchases of products/services is applicable towards







1 effective, apparently depending on the capriciousness and the judgment of the particular  
2 clerk or medical personnel in the carrier's office. These situations have produced enormous  
3 pressures to improve communications by extracting detailed data from providers pertaining  
4 to carriers failing to support appropriate and necessary health care delivery and  
5 documentation of how these insurance issues have caused definitive health problems. This  
6 data will be used to exert pressure on insurance carriers. In view of these major economic  
7 pressures, hospitals and other providers of medical services will definitely be most  
8 cooperative in providing detailed billing and clinical information with appropriate release  
9 consents, or alternatively through deployment and utilization of the proprietary techniques  
10 for privatizing patients' records using a pseudonymous proxy server. This information will  
11 be provided by the billing and medical records departments of the hospital and also, in the  
12 case of hospital employees, from the Human Resources Department of the hospital for  
13 reimbursement for medical services for hospital personnel. The minimal information should  
14 include the billing record for those in-hospital days or the specific services/tests denied and  
15 the discharge summary of that hospitalization. Based on reviews of this material, further  
16 portions of the in-hospital clinical record may be requested. (Additionally, the iamworthit  
17 service provides for the utilization of the above mentioned proxy server for collecting and  
18 maintaining pseudonymous patient data whereby advertising revenues from industry  
19 advertisers accessing/targeting desired user profiles may be shared with the hospital and  
20 potentially also the patient, thus further insisting the hospital to provide this needed patient  
21 data).

22 A compounding problem is the lack of information provided prospectively by insurance  
23 companies to hospitals, physicians and other providers. Subscribers and providers have  
24 the right to know, in detail, the number and kinds of services covered, and the depth of  
25 the insurance coverage including the length of stay. This information is legally  
26 accessible and can be made available not only to the patient, but also to his/her physician.  
27 On exemplary commercial need for the present service by employers who wish to  
28 reinsure their employees with quality coverage.

29 *Solution:*

30 In order to protect the consumer, a solution proposed to address a timely and  
31 overwhelming need on the part of consumers to establish a control data bank,

accessible to the consumer, which reveals general and specific instances of miscarriages by health insurance carriers of their fiduciary obligations to provide timely payment for essential health care services. Based upon data regarding the insurance company's stated policy and statistical analysis of the ultimately approved or disapproved coverage, performance criteria may be determined pertaining to the percentage likelihood of the insurance company covering certain services and length of hospital stays for future services to patients. In order to provide a robust and comprehensive statistical analysis, it is also required that a certain core sample of non-coverage denied service be collected in order to determine a variety of correlation criteria which affect the probability of coverage in addition to the identity of the carrier, also the type of prescribed medical care for associated illness wherein coverage for hospital stay or treatment is denied. Other correlations are possible, e.g., the medical history of the patient, the family history, etc. Of additional importance, statistics may also be provided which indicate the degree of adherence (integrity) of an insurer to provide coverage and to provide it with the level of depth of coverage which was purported by the insurer according to its standard claims policy. The iamworthit insurance service is accordingly designed with two primary case objectives:

- 1) Provide detailed statistics based on robust historical data regarding realistic expectations which the patient can anticipate regarding acceptance, denial and quality on the basis of various types of coverage given the context of the insurance needs and requirements of the buyer and based upon the specific services and depth of coverage purported by the insurer and any associated medical information regarding the buyer. The service could even anticipate, by review of the collective medical statistics the most likely type of coverage the buyer may need, based on his/her medical record specifically within the context of the associated determined probability of coverage and quality thereof for each insurer in view of this medical history of the buyer. In addition, iamworthit may provide a general informational resource where buyers can access such statistical information by carrier. Data mining tools may enable the buyer to extrapolate the desired statistical correlations as to his/her type of profile of medical or specific features of his/her medical condition or specific type of coverage

1 and thus in light of this robust data better equipping the buyer to determine which  
2 insurance provider and associated policy is most appropriate for his/her specific  
3 health care coverage needs;

- 4 2) This information may be accessed in real time by the buyer (or for e.g. buyers,  
5 employer) while browsing specific insurers' sites, receiving specific offers by  
6 insurers (typically requiring certain medical information) or browsing an insurance  
7 retail portal. Typically, buyers of the present service will combine their use of the  
8 present browsing-based information with the Time-of-Purchase Vendor Competition  
9 Services (described below). Additionally, a rating feature may also be provided  
10 which enables the buyer to rate (and subsequent buyers to observe the averaged  
11 ratings) the various vendors by various types of criteria relevant to the buyer's  
12 experience in receiving acceptance, quality and depth of coverage (as well as how  
13 well the insurer adhered to promises for coverage as stated in its standard policy for  
14 coverage), of course, this feature could be used if for insurance products on-line for  
15 later purchase off-line.

16  
17 Iamworthit's time-of-purchase vendor competition application to medical insurance (as  
18 applies similarly to numerous other e-commerce commercial domains) involves the  
19 following protocol:

20 The iamworthit enabled consumer provides specific needs or requirements as to what types  
21 of medical insurance products or areas of coverage are desired. These requests may be  
22 submitted either to a portal (or "mall") which connects the buyer to an extensive resource of  
23 medical insurance providers, or, alternatively, these requests may be submitted to  
24 iamworthit which provides its own default portal to these resources (which likely exists and  
25 is marketed as an independent marketplace for insurance purchasers on its own behalf).  
26 Upon submitting these buyer requests, iamworthit is able to provide numerous competitive  
27 insurance providers offering similar insurance products which in this application may be  
28 accompanied by certain credentials (such as the buyer's age and many other clinical  
29 parameters).

30 Additionally, the pseudonymized digital medical records of the buyer (which may also be  
31 requested by an insurer, typically in addition or instead of submitting a questionnaire



1 Thus, optimal offers are provided to the buyer by a "buyer agent". Additionally, the present  
2 service may with some modification be provided to buyers dialing into an 800 number.  
3 Typically, in this case, however, (as is also the case in the on-line implementation), if the  
4 buyer does not have access to his/her pseudonymous medical records, he/she would  
5 typically answer a questionnaire for each provider (which would typically consist of an  
6 aggregate questionnaire satisfying the requests of all relevant carriers which could then be  
7 submitted automatically, or an agent could be programmed to automatically complete the  
8 questionnaire as needed.

9  
10 *Commercial Alliances with Regional or National Health Care Organizations –*

11 Because of the tremendous restrictions in providing insurance coverage to organizations  
12 which provide health care services, it will be greatly to their advantage for patients and  
13 physicians to be able to benefit from more fair and competitive insurance coverage by using  
14 time of purchase competition and the associated coverage probability service (providing full  
15 disclosure of non-coverage and inequities in delivery of coverage). The primary joint  
16 marketing entities which will benefit most significantly from this service and thus will be  
17 the strongest commercial allies, include independent physicians, medical clinics, large  
18 physician groups, pharmacies and perhaps even pharmaceutical companies, providers of  
19 medical supplies, hospitals and home health providers. Of the above entities listed,  
20 physician groups, clinics, and hospitals do a tremendous amount of advertising directly to  
21 consumers (and this is becoming increasingly true via the internet for physicians in private  
22 practice), particularly through the direct advertising medium of the internet (e.g., advertising  
23 on their web sites as well as email). These organizations, in light of their direct marketing  
24 exposure to the consumer which they provide, would likely become powerful allies in  
25 assisting the promotion and exposure of the competitive insurance with full disclosure of  
26 claim information service to their consumers.

27 Secondly, health care organizations which market to health care professionals, such as  
28 physicians and nurses, including pharmaceutical companies, pharmacies, and home health  
29 care providers may indirectly benefit by informing their target customers about the above  
30 consumer oriented service.

1    *Quality Ratings of Medical Suppliers*

2    iamworthit's insurance application provides a natural entry into a second commercial  
3    application of the service. In a fashion similar to that of Time-of-Purchase Vendor  
4    Competition Scheme, for purchasers of insurance products, another useful commercial  
5    application of the present service is facilitating the time of purchase vendor competition  
6    between medical suppliers and equipment providers. The synergy by which this second  
7    commercial application is enhanced is as follows: physicians and hospitals will likely be  
8    cooperative allies to the service in providing patient information substantially as required as  
9    well as promoting (along with home health care providers) particularly the iamworthit  
10   insurer coverage probability service to their consumers. They will also likely utilize such a  
11   resource for internal uses and if/when HMO service liability legislation becomes enacted,  
12   for legal purposes as well.

13   To the buyer's advantage, it is possible (as in the insurance application, albeit less likely)  
14   that unfair or arbitrary price discrimination may be detected and counteracted by  
15   iamworthit's data sharing feature. Common buyers in this environment include hospitals,  
16   physician clinics, home health care providers and to a lesser extent HMOs and end-buyers,  
17   the nature/types of supplies and equipment which tend to be purchased by each of the above  
18   categories of buyers are provided in further detail below. As suggested earlier in this  
19   description (and in issued patent "Pseudonymous Server for System for Customized  
20   Electronic Identification for Desirable Objects"), a vertical portal as the one described, may  
21   be personalized based upon the user profile as disclosed upon accessing that portal in  
22   general, at the item level as within a category (or in conjunction with search results).  
23   Additionally, vendors may be equipped with price discrimination tools and are in fact highly  
24   motivated to form a coalition by entrusting their data and counter offer responses to  
25   incoming time of purchase vendor competition offers and bids to SDI (or at least an industry  
26   specific counterpart of SDI) for medical suppliers in order to provide limits or "price  
27   ceilings" for any given buyer and associated profile and condition relating to competitive  
28   offers or bids.

29   As in the case of almost any category of purchaser, particularly the larger entities, such as,  
30   unions, hospital systems and large physician organizations, additional bargaining leverage  
31   can be applied to suppliers because of the volume of supplies which can be purchased. The





1 iamworthit medical supplies application as it similarly could apply to all of the above  
2 entities.

3 (e). Home Health Care Market - Equipment and supplies which are, in part, similar to large  
4 physicians clinics as well as some hospital - type supplies, used to provide out-patient care  
5 and maintenance services.

6 *Patient Medical Records*

7 The secure data interchange provides a secure and privacy protected storage and retrieval  
8 architecture in which buyers (in this case patients) are able to control the use of their  
9 personal profiles (in this case medical record information). Herein, we propose an  
10 alternative model which is based upon the assumption and belief that purveyors of this  
11 medical information, including hospitals, physician's offices/clinics, labor unions, (to some  
12 extent) HMOs are legally permitted to disclose and use their information for advertising  
13 purposes so long as a method is provided which insures that their information will not be  
14 associated with any identifying information about the patient (name, address, social security  
15 number, etc.). If regulatory constraints and/or the purveyor of this information prohibits the  
16 use of this information for the above purpose, patient permission could be acquired by  
17 patient signature on all appropriate consent forms at the time of admission and directly on  
18 the admission form.

19 In one variation, software may be provided to the organization which enables the  
20 organization to operate their own pseudonymous proxy server, thus enabling access by  
21 advertisers to the pseudonymized patient records. Email and/or telephony-based  
22 pseudonymous communications could be readily provided. For direct mail, the address  
23 information maintained by the hospital would mandate the hospital to provide direct mail  
24 advertising to the patients or to entrust this operation to a pseudonymous physical mail  
25 service such as SDI. In the (much) preferred variation, a trusted pseudonymous proxy  
26 server operator maintains the patient information on an external database (this would  
27 perhaps be SDI or an SDI-like service) in which all of the above pseudonymous  
28 communication media could be provided along with pseudonymous direct mail services.  
29 For the SDI data entrusted variation, the preferred business model to provide incentives  
30 to these organizations to cooperatively provide this data would involve SDI's privacy  
31 policy enforcement capabilities in combination with a revenue sharing model in which

1 revenues from advertisers such as pharmaceutical companies, pharmacies, clinics,  
2 HMOs, hospitals (where the purveyor of the data is not a provider of the same  
3 competitive services). In the case in which HMOs are the advertisers, the HMO would  
4 identify patients who are currently subscribed to competing HMOs and (typically) have a  
5 medical history indicative of a good risk to the HMO. The iamworthit insurance  
6 competition scheme, as above described, could be implemented within the context of this  
7 current solicitation scheme in order to elicit time of purchase competition among HMOs  
8 upon the submission of an offer by the original HMO to the pseudonymous patient.

9  
10 advertising purposes so long as a method is provided which insures that their information  
11 will not be associated with any identifying information about the patient (name, address,  
12 social security number, etc.). If regulatory constraints and/or the purveyor of this  
13 information prohibits the use of this information for the above purpose, patient permission  
14 could be acquired by patient signature on all appropriate consent forms at the time of  
15 admission and directly on the admission form.

16 In one variation, software may be provided to the organization which enables the  
17 organization to operate their own pseudonymous proxy server, thus enabling access by  
18 advertisers to the pseudonymized patient records. Email and/or telephony-based  
19 pseudonymous communications could be readily provided. For direct mail, the address  
20 information maintained by the hospital would mandate the hospital to provide direct mail  
21 advertising to the patients or to entrust this operation to a pseudonymous physical mail  
22 service such as SDI. In the (much) preferred variation, a trusted pseudonymous proxy  
23 server operator maintains the patient information on an external database (this would  
24 perhaps be SDI or an SDI-like service) in which all of the above pseudonymous  
25 communication media could be provided along with pseudonymous direct mail services.  
26 For the SDI data entrusted variation, the preferred business model to provide incentives to  
27 these organizations to cooperatively provide this data would involve SDI's privacy policy  
28 enforcement capabilities in combination with a revenue sharing model in which revenues  
29 from advertisers such as pharmaceutical companies, pharmacies, clinics, HMOs, hospitals  
30 (where the purveyor of the data is not a provider of the same competitive services). In the  
31 case in which HMOs are the advertisers, the HMO would identify patients who are currently

1 subscribed to competing HMOs and (typically) have a medical history indicative of a good  
2 risk to the HMO.

3 The iamworthit insurance competition scheme, as above described, could be implemented  
4 within the context of this current solicitation scheme in order to elicit time of purchase  
5 competition among HMOs upon the submission of an offer by the original HMO to the  
6 pseudonymous patient.

## 7 2.2 Matchmaking 8

9 The user-centric SDI model allows users to provide personal information on a carefully  
10 controlled basis to vendors and other users. Furthermore, vendors can implement rules that  
11 personalize the information, products, and service provided to users—on the basis of  
12 personal information that they receive from users directly, or have acquired about users. The  
13 key property of “matchmaking” applications is that information exchange is bi-directional,  
14 i.e. a match is not computed on the basis of static information within the system, but may  
15 involve additional information exchange between agents, and/or special actions by one or  
16 more agents.

17 Client-level SDI proxies can act as autonomous agents in an architectural variation of SDI,  
18 where the “client-level proxy” is co-located with a (physically) mobile user, for example on  
19 a palm-held computer or head-up display. In a “match-making” application the goal of these  
20 SDI proxy agents is to find other agents with desired profiles/synergies, with the view to  
21 taking an action or exchanging a particular type of information. We describe many  
22 applications, for example: to matching buyers with a virtual sales force, to a market for  
23 expert opinion, to groupware systems for virtual team working, and to an investment  
24 system.

25 We can allow agents to communicate *anonymously, without revealing (even*  
26 *pseudonymous) identities*. In addition, we can prevent one agent receiving profile  
27 information about another agent unless authorization is provided within the rules of data  
28 release associated with the agent’s profile. Implicit authorization occurs when a requestor  
29 can present certificates to verify that it has required attributes to access particular  
30 information. Explicit authorization occurs when a requestee provides direct authorization  
31 to a particular part of a user’s pseudonym.

1  
2 Essentially there is bi-directional information filtering: the requestor agent will only present  
3 certain information to the user, information that is relevant; and the requestee will only  
4 provide information when a request is judged to be legitimate. Information exchange  
5 between agents occurs as part of a multi-step negotiation, until both parties can agree on  
6 terms for either a physical meeting (or execution of a deal), or further pseudonymous  
7 exchange of information or cooperation.

8 In match-making applications the role of the central SDI query execution engine is to search  
9 continuously for "matches" between agents, based on agent profiles. When a match occurs  
10 additional information exchange may occur automatically between agents, until finally an  
11 introduction is made or an action executed by one (or both) of the agents. Decisions about  
12 what information to exchange are made on the basis of both static and dynamic profile  
13 attributes, e.g. standard (historic) profile information, current behavior, current location, and  
14 recent activity. Also mode of behavior e.g., social, business, leisure as detected by the  
15 user's calendar,, time, content in recent conversations, type of users recently interacted with,  
16 etc. LEIA style-behavior attributes can be used to automatically decide on the relevance of  
17 new virtual tag information. A requestee might also demand certain credentials to indicate  
18 the lack of negative reputation marks, for example that an interaction with the user has never  
19 received a bad rating (see Section 2.2.3). Perhaps a third-party could be used to determine  
20 whether the user's know each other (e.g. [www.sixdegrees.com](http://www.sixdegrees.com)). When a requestee denies a  
21 request for information, it may instead provide criteria for data releases. A requestor can  
22 respond with a different information request, or a subset of required credentials. Finally, the  
23 agents might agree on terms of negotiation and conditions can be anonymously fixed.

24 There are (at least 5 levels) of information disclosure: indicate to another user interest;  
25 release profile information; disclose communication; start a correspondence session;  
26 schedule a meeting/strike a deal, take some other action, etc. The end-result of information  
27 exchange could be an agreement to calendar a meeting for some future time and place; and  
28 absolute, or pseudonymous revelation of identity.

29 An initial implementation of the data-release policies might allow only manual definitions.  
30 However, after an initial "beta testing" phase, a data mining suite could be used to cluster  
31 users and generate exemplar data release and data request policies. A system can provide

1 default settings for users, and recommend setting based on users with similar profiles. The  
2 user can further fine-tune the rules. Automatic feedback techniques can be useful to adjust  
3 rules, for example—when a user is especially receptive to particular type of introduction  
4 then make such introductions more likely in the future. An intelligent interface system might  
5 also suggest refinements to the rules, to automatically cover “patches” where the user  
6 currently controls interactions. For example, some rules may be based upon a certain  
7 confidence threshold, which the system suggests or refines. The user may accept or adjust  
8 the threshold as desired. S/he may wish to provide settings as to which types of autonomous  
9 actions the user wishes to authorize updates to (based upon his/her own behavior or as a  
10 default users who are very similar) and the user wishes to authorize the recommended  
11 action, or allow the system to fully automate the activation of the recommended actions (or  
12 do so only if the confidence threshold is above a certain level). Or for example bypass user  
13 authorization of a recommended action if it falls within a certain margin of confidence  
14 below the normal confidence threshold for activating full automation if the user is  
15 performing another important task , e.g., in a meeting, engaged in an important phone  
16 conversation

### 17 **2.2.1 Resolution Credentials**

18 Resolution credentials from third parties can play an important role in agent-agent  
19 interactions. We can prove the absence of a quality attribute or behavior (which is often of a  
20 negative nature) relating to an individual and is submitted by a third party and typically must  
21 be issued on a periodic basis in order to maintain currency. A few simple examples of  
22 resolution credentials which may be of interest to users (credentials which users may  
23 commonly request as a precondition to requesting or accepting requests to be introduced or  
24 initiate communication with an outside unknown third party) include:

- 25 1) For business associations: are in good business standing, e.g., have not attempted  
26 to defraud other users in the course of common business practices. Or maintain  
27 sufficient funds in one’s account to perform business activities (as represented by  
28 the user).
- 29 2) For business interactions or social interactions: are in good standing with the law.
- 30 3) For social interactions: have not been accused by other individuals of  
31 inappropriate or antisocial behavior.

1  
2 Some standard credentials which may be of interest to many users, and which may (as with  
3 resolution credentials) be incorporated with the standard settings of the user's data request  
4 policy as herein described. A few examples are cited (among countless potential others):  
5 profession, awards, honors, alma mater, e. g., Harvard graduate, doctorate degree, etc. In  
6 accordance with the parent issued patent US Patent #5754938, various credential issuers are  
7 provided for issuing standard and resolution credentials to individuals. Thus certain entities  
8 may be entrusted with "legitimate authority" to validate and submit credentials which are  
9 issued to the appropriate individuals. If a resolution credential is not issued (or not renewed)  
10 an adjudicating third party is provided which has access rights to both of the parties is  
11 provided to resolve resulting disputes (from the subject user). The present invention  
12 describes how credentials can be issued to users pseudonymously.

13 There are a variety of rules which a user's data disclosure policy and data request policy  
14 may contain, to control what if any attributes are released, and what credentials are  
15 required. A data request policy may state a rule for explicitly notifying the user if a  
16 particular resolution credential (e.g., indicative of a serious problem or concern) cannot  
17 be presented in response to the user's disclosure request. We allow initial information  
18 exchange to be anonymous, such that information that is released as preconditions for  
19 release of further information is not useful. Similarly, so long as initial encounters are  
20 anonymous there is no need to withhold information about them from the user.

21  
22 Some users may not wish to disclose specific information about themselves via these  
23 standard credentials but instead certain "extracted" more general information may be  
24 provided about themselves. For example, instead of a "Harvard grad or Ph.D." there may be,  
25 for example, credentials indicating "intellectual" or "prominent intellectual". Or instead of  
26 indicating an individual's wealth or value of assets, the credential may indicate "wealthy" or  
27 "very wealthy" (typically, depending upon user's wishes this latter credential should also be  
28 withheld during initial introductions or subject to some fairly stringent conditional criteria  
29 from the other party) and instead replaced with an even more general credential e.g.,  
30 "prominent" or "influential citizen"). Similarly, an individual's exact profession or scope of  
31 work may not be fully disclosed initially but rather a more general definition of his/her

profession or perhaps the general field initially in which the user works or his/her particular views on certain social or political issues may instead be substituted with (for example) “socially conscious”, “philanthropic”, or “moderately liberal”.

In accordance with the parent patent application, rules may be learned regarding certain things that a user does (as in ascribing these rules for which messages to send to whom or what user profiles and under what circumstances/events surrounding the target user). Thus, his/her agent may begin to suggest certain future actions which could be performed in the future upon user approval or even automatically. If the user has had no previous interaction at all with the system, it may identify which other users of the system the present user is most similar, and recommend initial rules. Additional textual attributes can also be leveraged to provide extra criteria, and data mining techniques used to generate more appropriate rules.

Additionally, the communications which the user may be presently involved in i.e., the content profile of his/her spoken dialogue and/or other “on line communications” may be used and combined with location/time patterns in order to further infer the circumstances,



1 behavior, and present temporal interest of a user and/or third party for purposes of  
2 employing the user's data disclosure and data request policies.

3 Credentials can allow users to identify other users that may pose a threat. This identification  
4 may be provided vis-a-vie resolution credentials and/or rating (by third parties). e.g. a user  
5 has not engaged in any serious criminal activity, physically harmed another person, or  
6 interacted with other individuals who are unable to produce these resolution credentials.  
7 Other credentials may specify the nature of an infringement, and its context and severity  
8 (e.g. what was the context of a physical assault? Was it performed during a bar brawl,  
9 against a friend, a boss, an elderly person, a child, a family member – or at work? In this  
10 case, the user agent may, for example, bring to the attention of a prospective employer that  
11 the user could not present a credential indicating that they had not previously harmed or  
12 threatened a former employer. Was it minor or severe? Also, if such individuals (lacking,  
13 for example, resolution credential proving the absence of having committed armed robbery)  
14 are (or come) within a certain proximity of a user, the user may wish to program his/her user  
15 agent to notify the user. The same would, of course, apply to a store clerk regarding  
16 customers of this sort or to baggage security personnel at an airport. Or, highway patrollers  
17 may be interested (e.g., on certain stretches of highway) in being made aware of vehicles  
18 and their locations whose agents are unable to provide a resolution credential proving the  
19 absence of a drug conviction.

20 In another application (in accordance with the auto insurance risk determination methods  
21 described in co-pending patent application entitled "Applications for Location Enhanced  
22 Information Architecture" [INSERT PATENT OFFICE NO]), an on-board computing  
23 device within a user's automobile could identify another automobile lacking, for example, a  
24 resolution credential for safe driving. i.e. the on-board user agent continuously polls agents  
25 in other cars for a "safe driving" credential, and if it fails to receive such a credential it  
26 issues a warning to the user. As an extension, this location data could be converted into a  
27 dynamic 2-D rendering upon the user's windshield (using heads up display technology) in  
28 order to thus superimpose a persistent flagging or highlighting of that particular automobile  
29 from the driver's visual perspective. Pedestrians or law enforcement officials (for example)  
30 could also receive instant notification. As is described in LEIA, a roaming cellular  
31 connection, or GPS, is not essential for providing a user identifier. For example, optically-







1  
2 In light of the present context of the situation and potential opportunities, the questions  
3 may be selectively presented in order of priority and are typically related to the attributes  
4 characterizing present opportunities/individuals which are either of predicted interest or  
5 of unknown predicted interest/relevance but may possibly also include potentially  
6 predicted relevant questions about the present frame of mind of the user, such as what the  
7 user's present activities are (i .e., including work , leisure, etc.) mood, present focus on  
8 what kind(s) of content, etc.

9 In the above example, of course, Mr. A and Mr. B could instead be established friends or  
10 colleagues in which case keeping secret their present physical locations may be  
11 considerably less important but not necessarily so. Revealing informational details about  
12 location, action, attention focus to another third party while completely leveraging this  
13 information to provide optimal scheduling efficiency is all possible with SDI and many  
14 individuals may wish to not disclose it in this way (as a default) even if the disclosees are  
15 relatively trustworthy.

16  
17  
18 Complete robust functionality coupled with complete user privacy as provided above  
19 through the user's agent provides a compelling motivation for users to maintain much of  
20 the information, regarding their present and general location and activity related data,  
21 private in the majority of cases of user interaction and associated schedule coordination.  
22 If Mr. B were to pose different relevant queries to Mr. A's agent while determining  
23 whether he could be interested in meeting, for example, where/when their paths may  
24 cross based upon different times and/or points of departure, the specific location and  
25 travel information regarding Mr. A. may still be fuzzed as the meeting time at the  
26 associated place of rendezvous does not disclose to the other party when the other will  
27 arrive, only that time when both will be there at the same time(for those rare cases where  
28 even this information may disclose likely specific location and intended time of  
29 departure, we can apply randomization techniques in order to assure privacy with  
30 relatively minor negative impact upon scheduling efficiency. Nevertheless, how much  
31 privacy versus efficiency is desired is entirely within the control of the user.

In another application of the system, the individual may wish to allow another user (or entity) to be able to arrive at certain levels of deductions or conclusions about the user which may concern them without explicitly revealing any facts or details used to arrive at those conclusions. Or in a variation, even allow a continuous persistent revelation of information regarding these certain conclusions such that if a certain conclusion(s) occurs (or ceases to occur) notify the requester of such facts. A statistics and/or manual rules-based approach could be used in this instance. A statistical interface provided to SDIs central data warehouse could enable an expert to establish correlations and confidence thresholds appropriate for deducing certain conclusions. For more novel or complex rules which haven't been seen before SDI could receive disclosure from users regarding their actions and a request as to desired types of conclusions. The industry could further trigger an anonymous action (and/or notification) on behalf of the user or depending on the privacy policy (consent of the other party) an action could even be triggered based upon such conclusion in the absence of disclosure of the conclusion and/or even the autonomous action taken.

E.G. notify the subject user of interest in message (or suggestion or warning) another relevant, noting third party, etc. In such cases, it may however be acceptable for the user to be able to receive fuzzed aggregate statistics regarding such data as a user's past behavior patterns (or more confidentially) other users who share similarities to the user regarding his/her user profile and/or perhaps context of similar activities and circumstances. User profile could be more general characteristics of the user than the detailed user profile as suggested herein, e.g., the user's corporate department peers, class, demographics, psychographics, travel venue or combination thereof. This process may be performed in either a pull or even push mode. Many associated applications are conceivable, for example, a spouse (the requester) may desire to have his/her agent persistently issued a resolution credential indicating that all of the activities of the user while on a business trip are consistent with behavior which is appropriate for a married individual, e.g., that the associations with business associates/colleagues are maintained at a professional level (per that mutual physical locations and perhaps even directly or transmitted spoken or written communications) or that user doesn't engage in late night carousing at questionable venues, or , for example, that the individual on his way to

work, to a meeting worked diligently on an important proposal or on his/her way home from work without revealing any further details which his/her agent has explicit knowledge of with which it deduces these facts with confidence. Another example might involve the use of LEIA via small cellular transmitters (with or without GPS) affixed to small children or the elderly whereby location and activity related information is maintained private, unless an inappropriate action is performed (e.g., motion and/or acoustic detectors in the elderly person's home suggest s/he may have fallen, that the child has wandered away from parents or is conversing with an inappropriate stranger. In such cases, notify the guardian or assistant in charge or allow a highly trusted party who happens to be in the immediate vicinity of the party to be notified to come to his/her aid. In the latter case (of the child) perhaps there is suspicion that the child has just disappeared and the parent(s) authorizing the agent of the child is able to send out an immediate description of the child to all certified "trustworthy" parties who are in the immediate vicinity, e.g., after the initial description has been sent out, one of the parties indicates the child was speaking or walking with an apparent stranger of X description which is also, in turn, transmitted to these trusted local parties (and at that point to the police). In a variation of the above example, a crime could have just been committed and the assailant has just been described in detail by the victims with information regarding location and direction last headed which is again transmitted to other parties: Thus, more immediate, complete and up to date information of importance for police to make a more immediate apprehension of the suspect is made available to police.

**INSERT>>>>>>>>>>**

Location determination via LEIA could be either through user's LEIA-enabled devices or the coordinates of the user(s) as identified verbally by the user, e.g., to a 911 server.

Other examples of the present system could apply to immediately locating trained medical, EMS trained or rescue personnel in proximity of a urgent medical emergency or accident. Such personnel could even be off-duty. Depending upon the particular nature of a medical emergency, particular specialists appropriate to the circumstances could be identified, their comparative physical proximities compared with other potential qualified individuals and based upon qualification appropriateness, physical distance and perhaps willingness, a most appropriate match made. In another application, likely witnesses to

1 an accident (including automobile accident), theft abduction based upon their location  
2 and time relation to the event of interest could be notified to provide helpful clues in  
3 assisting investigators in obtaining a suspect or legal professionals a conviction or  
4 acquittal. In the case of abductions timing in acquisition of such information could be  
5 extremely critical and the ability of such notification to be distributed instantly to all  
6 "trusted" individuals in the immediate vicinity at the very outset of when a person (e.g.,  
7 child) is even suspected of being lost could prove to be invaluable in not only rapid  
8 recovery but also in thwarting of those few instances in which a kidnapping has actually  
9 occurred. Some users may wish to be available and willingly notified to help in certain  
10 dire circumstances (the nature of which the user ultimately has control over) but may  
11 wish to remain anonymous to authorities and legal professional after the fact (or for  
12 example in cases of providing witness testimony (as via a subpoena) to more trivial or  
13 non-life threatening cases. A market model may also be created in which compensation  
14 schemes which are just appropriate to compel typical individuals to provide desired  
15 assistance may be provided. Finally, the present methodology could be readily extended  
16 to a taxi dispatching service in which a taxi's most local to a given user can be dispatched  
17 accordingly. Another example includes employees within a large organization. The  
18 statistical techniques as described in the parent patent application could be trained such  
19 that examples can be identified or explicitly stated rules provided by employers of what is  
20 considered to be actions and behaviors which are inappropriate and irrelevant to the type  
21 of work related activities of that employee's job description (ideally trained across  
22 previous "similar" employees). The system can take into account browsing,  
23 communications, (including spoken communications), even location (LEIA) data into the  
24 model (e.g., for tele-commuters) and flag suspicious behavior. SDI can even upon  
25 agreement by the employee) allow the employer to access more detailed information on  
26 the specific behavior once the suspicion threshold has been exceeded.

27 In another application, the user may wish to provide predefined rule(which could be  
28 suggested by the system initially in (similar fashions as suggested above).s which can be  
29 used to notify a certain individual(s) if certain explicit (or agent) inferred events occur,  
30 for example, that the user is in town, has just received certain important news of a  
31 promotion or change to a new position, internal company news which relates to that



individual and may have professional relevance on the other individual (the latter two examples which may be relevant to say the professional profile of that individual , etc.). In addition, the individual(s) may be made aware of the event disclosure policies set forth by the user relating to them specifically (if desired) and this individual(s) can further provide filtering rules which allow all or a portion of that information to actually reach the user. In another example, the user may have met or known the individual at some time previously and the venue and context of the meeting circumstances and relevant user profile information to that context are known and disclosed by both parties. The user then wishes to re-establish contact perhaps defining the reason for the request and the individual can have predefined rules or acceptance or rejection of the request directly. In a final example, virtual tags are provided containing comments and annotations regarding a certain physical object (which could be anything from a tourist attraction to the physical site of a recent hot news event to a useful annotation critiquing a bricks and mortar vendor). An individual whose profile suggests s/he is knowledgeable (or had provided the relevant annotation of interest) about that object would allow this relevant information to be conveyed to other individuals who have perhaps an interest in that information (generally or at that moment).

The relevant individual (if s/he is also willing) may be contacted and solicited to meet if their calendaring agents can appropriately coordinate the meeting. It may be either context specific, e.g., both individuals are physically local (or plan to be local to each other) and the physical object or site of interest or simply (generally) local to each other.

It is perhaps worthy to note that the present methodology could certainly be extended to include the application to rules dictating reachability conditions of users bu other users via standard telephone. The application to cellular telephony is of particular interest in its ability to communicate to the user in a variety of modalities (voice, notification, instant messaging, Web content). In addition, similar rules may be adapted in this scenario with perhaps minor modifications. For example, under what conditions may the user be reached by another user. E.g., what is the nature of the relationship between those individuals? If it is not known, inferences from the context of previous conversations may be determined. Does the user typically accept calls (or other communications) from the other user under

1 similar circumstances and/or times. If this data is not available are they accepted by other  
2 users similar to that user according to a key attribute(s) or does the other user accept  
3 communications from the user under similar circumstances as they present communication  
4 attempt (as this may suggest similar reciprocal acceptable reachability conditions as that of  
5 the other user. Again, LEIA suggesting location movements and associated measurable  
6 behaviors as sequential patterns and as formation of time may be important. This detailed  
7 statistical model of the user could be uiseful in learning the context of performing certain  
8 autonomous actions for the benefit and convenience of the user and in the reachability of the  
9 user by other users, e.g., in what communications mode is the user interested in  
10 communicating, e.g., speech to text-based instant messaging or standard telephony, does the  
11 user mind being interrupted with phone ringers, (or vibration mode) or ringer-off mode).  
12 Previous similar communications contexts may be useful in determining inferred rules  
13 which may be presently appropriate (e.g., was the spoken content of the previous  
14 communication indicative of acceptance or annoyance of the communication with the other  
15 user within a similar associated context . Eg., this may suggest whether and in what  
16 contextual conditions of both accessor and accessee the accessor's call may be patched  
17 through (or re-forwarded if initially accessing another number or telephone of another likely  
18 venue where the user may be believed to be present.

19  
20 Exemplifications of The Underlying Statistical and Rule-Based Intelligence Used in  
21 Autonomous and Semi-Autonomous Release of Personal Data Release, Matching and  
22 Reachability by Other Users  
23

24 The above section "Location Enhanced SDI System relating to Smart Home and Office  
25 techniques provides a relevant platform for integrating intelligence into a multi-device  
26 environment. Mobile user intelligence is a sub-component of this broader multi-device  
27 (or "ubiquitous computing") application-level intelligence platform. In view of the  
28 present application framework, it is a very intriguing and challenging problem to provide  
29 a statistical framework which is able to allow the user agent to make appropriate  
30 inferences regarding users who generally, though imprecisely, fit a certain profile which  
31 may suggest the activation of a rule (i.e., a "fuzzy rule") such as requesting further





1 locations or mind sets experienced by the user. E.g., a young female in her twenties may  
2 not wish to be approached by 40 year old males unless the context of the interaction were  
3 purely business, e.g., a professional, or commercial or sales opportunity thus if she were  
4 in a purely social mindset such requests to interact would be denied. Likewise a busy  
5 executive in his/her 30's may deny such requests from the young female unless her  
6 request was purely of a social nature and he was not overly time constrained. There may  
7 be certain "professionally" prospective correspondents, however, which he may prefer  
8 over others based upon some additional interesting criteria as well such as other  
9 business/professional areas of commonality, common interest areas, even socially  
10 compatible or interest criteria.

11 Of course, SDI is able to extrapolate extremely sensitive features regarding interest  
12 preference and context which may establish a basis where two individuals or more  
13 individuals, to be mutually identified to one another and/or communicate accordingly or  
14 to not be mutually identified and/or accordingly communicate (SDI may even establish  
15 the basis for these mutually beneficial interactions even if revelation of this basis by one  
16 or more parties by the other is not permitted within the privacy policy of one or both of  
17 the parties). Of course, SDI may frequently identify individuals whose personal data  
18 revelation requires certain minimal requested personal data from the other party as a  
19 condition to such revelation (which SDI is capable of fully automating on behalf of both  
20 parties) or the user may request certain information and make a personal judgment in real  
21 time as to whether further information about him/herself should be revealed and/or  
22 whether introduction or communication is desirable. In these cases it may be difficult to  
23 always successfully achieve an introduction when appropriate as the user is unable to  
24 guarantee

25  
26 a). The revelation of certain information to the other use if that user reveals certain  
27 information accordingly, or

28  
29 b). Is willing to agree to be introduced if certain desired information is revealed by the  
30 other party (often busy executives have absolutely no time to disrupt their busy schedules  
31 to discuss certain opportunities when there is enough revealed about those opportunities



1 which is mutually compatible (similar applications are suggested for matching  
2 sales persons with prospective clients, identifying experts to work (individually or  
3 collaboratively) on a particular project or problem, to answer a question of an  
4 appropriate specialized nature to their area of expert knowledge.) The parent  
5 issued patent suggests at a general level these commercial applications. An  
6 additional feature described therein involves the use of a decision tree called  
7 "Rapid profiling" which can be used in the present context to identify from the  
8 most common needs of users and "goods" of sellers in general and the known  
9 profile data about each user and seller individually, a list of questions for each  
10 party which most briefly and efficiently determines the complete user/seller  
11 profile of each party individually.

12  
13 Social Interests Profile Information—The parent issued patent also suggests the  
14 present application at a general level. For a dating application, users may be  
15 matched on the basis of their common interests/preferences and perhaps on the  
16 basis of certain information reflecting personality, social or cultural  
17 behavior/affinities or psychological attributes. On the other hand, for purposes of  
18 meeting casual acquaintances, users may be interested in another user who shares  
19 the above characteristics as well as someone who has recently shared similar  
20 experiences and/or personal challenges.

21  
22 Professional Information/Qualifications - As in the application of matching users and  
23 sellers, a description of a user's needs or situation with relation to various  
24 professional services may be provided as additional data about the user.  
25 Examples may include: (as above) medical data, professional or business history  
26 (as well as legal history) which may be of interest to law firms, accounting firms  
27 or various business consultants. Personal, family or emotional difficulties may be  
28 of interest to psychologists or family counselors. Again, users may submit this  
29 information as a query for prospective matches, or they may be pseudonymous  
30 queries or automatically matched in accordance with criteria specified by the  
31 professional. The issued parent patent application also lists additional

1 applications, which could as well be relevant within the usage context of virtual  
2 tags.

3  
4 Employer/Employee Information - An employer may post a description as part of  
5 his/her virtual tag (and that associated with his/her company). His/her employees  
6 may also have provided ratings and/or annotations, which are further descriptive  
7 of his/her personality, leadership/management style and skills, work environment  
8 which s/he promotes and overall quality. A previous employee may allow  
9 him/herself to be contacted by the prospective candidate (e.g., in exchange for a  
10 fee).

11  
12 Access Privileges Information - Users in an organization are frequently given  
13 privileged access to certain files within a corporate intranet but not others.  
14 Though there are many ways of profiling users according to their level of access  
15 privileges to information, the following example is considered: Based upon the  
16 position (e.g., responsibilities and tenure with the organization), users may be  
17 "classified" into groups according to different levels of access to confidential  
18 information. Virtual tags may be used to extend the capability by providing for  
19 immediate disclosure of a user's information access privileges to another  
20 employee in real-time and in a physical context. Also, if a user reads or accesses  
21 certain information, meets with a certain colleague or friend, then a user might  
22 send a message X. This message could be (for e.g.) a request to perform some  
23 task relating to part of that information, a reminder to address certain issue(s)  
24 while chatting with the colleague etc. or, per the request of an individual's  
25 employer or colleague if a given individual (a sales person) meets with user X  
26 send him/her message Y (which may refer to a previous encounter, experience or  
27 fact s/he should know pertaining to user X and which may have bearing upon  
28 their conversation or professional interaction.

29  
30 Access restricted physical areas. In this case the virtual tag effectively may behave  
31 like an "electronic door key". A variation of the technique may be used for





benefit from this information as well as provide information for other vendor identification of competition. Other data resources including electric payment protocols, EDI, automatic check payment, check services, etc. may be useful data resources as well.

Again the disclosure of detailed business information is very helpful and a data release policy defining the parameters for such strategic initiatives may be critical in order to determine what companies may be potential candidates for which initial feelers (of high level information disclosure) would be appropriate to put out to a prospective company to determine mutual interest and/or further basis for expected synergies.

#### 2.2.4 Dynamic Annotation/Information Filtering

In this extended application of SDI, we allow users and other third parties to annotate objects (physical and virtual) with meta-information, either to remind themselves about a previous interaction in the future—or as a system of “knowledge learning”, where systems of users leave useful information for other users. Information is left in the environment, leaving a trail for other users.

For example, the information that is tagged to an object, referred to as a “virtual tag”, can contain a pointer to other relevant information, such as a survey of a film by a third party, or the user’s own comments/feedback. For example, a restaurant listing could be annotated with meta-information about the quality of the food and service. Such information, when provided by a wide sample of users, can provide robust information about objects. The information that is used by a particular user can be filtered—for example, weighting the opinion of a respected restaurant critic, or weighting the opinion of users with common profiles (when that information is available).

Virtual tags (i.e. meta-information) can be assigned to objects with physical locations, and the information triggered based on the physical location of a user (using LEIA technology).

Virtual tags can be assigned with expiration dates or other time-sensitive information. An individual user might leave an “action item”, for example—next time I return to this object (e.g. web page/ vendor) be sure to perform this task, enter this query, check this link for new information. As another example, after a conversation with an SDI-enabled user it is possible to tag that user with some notes, to remember the conversation the next time the two users meet.

The technical innovation that allows this use of virtual tags, in addition to the protection of privacy, is that we allow users to annotate information to objects that they do not directly own through a system that separates virtual tags from the content that is tagged. In particular, tags can be stored (either at the ISP-level proxy, or main SDI server) for associated web pages, and exchanged/retrieved automatically when the object is accessed. The virtual tags can be used in conjunction with target-object profiles that are generated through SDI for web pages (and approved by vendors). Virtual tags can be searched, using relevant terms, locations, or times, and can also contain links to authoritative information, such as audio and/or video.

For example, in accordance with the prediction that readily deployable visual video recording devices will become commonplace (even integrated into wearable computing hardware). Users who are amenable to releasing such information under terms of their privacy policy may allow other individuals subsequently visiting the same physical location physically or virtually to gain access to such information. Such information may range from detailed accounts, assessments of value, etc. A user could, for example, do a general query about locations e.g., what is the percentage breakdown of types by their pseudonymous attributes, who frequently visit this site which sites tend to have commentary of a particular topical nature of interest to user U, which sites are visited frequently by individuals similar to user U, which sites have had an event of a particular type or one which is similar to user U's user profile occur in the recent past (or where or how close did such event occur with respect to the recent location of user U and so forth).

Tags are encrypted, so that only SDI-enabled users can access them. Tags are also associated with the pseudonymous ID of the user that left the information (although they can be anonymous, an associated profile allows more accurate collaborative filtering techniques). Finally, users can leave data-disclosure policies, embedded into tags—to certify the properties of other users necessary to release the information. When tags automatically are time-stamped with location, and time, and other information we allow for this information to be “fuzzed”, as disclosed in the section on Randomized Aggregates, to protect a user’s identity.

In the physical world, implementation of meta-information in a user's physical information, can be viewed via head-up displays, video cam monitors, wearable computing devices, or

1 audio pieces. The information itself can be embedded directly on physical objects, for  
2 example on magnetic strips or via visual encoding techniques—or the appropriate  
3 information can be accessed from a secure remote database based on the user's physical  
4 location (using LEIA location technology); or bar-codes that provide a universal identifier  
5 for an object.

6 As an extension to this model, we also allow users, vendors, and other third parties to  
7 associate “meta-information” with other *users and vendors*. This information might be a  
8 user's opinion about his/her interaction with another user, an annotation that relates to a  
9 particular web page, or information about a physical object. The system of SDI enhances the  
10 value of this information by providing a secure environment where users can report meta-  
11 information (i.e. opinions) along with their profile information, to permit *collaborative*  
12 *filtering techniques* to generate appropriate meta-information about an object (user, physical  
13 object, vendor, web page, etc.) that will be useful to a particular user—given that user's own  
14 profile. We define “virtual tags” as any piece of information about an object (physical or  
15 virtual). The information may be authored by any party, but annotated accordingly. For  
16 example, the appropriate virtual tag provided by a user about his/her-self is the  
17 pseudonymous profile for that user, -- and with SDI only the user his/her-self can gain  
18 access to the profile (either directly through editing, or indirectly through continuing  
19 transactions).

20 We might implement a Kasbah-style “reputation system” within such a virtual community.  
21 Initially users (under pseudonyms) have no reputation, and their opinion does not count for  
22 much, but after every positive interaction (as defined by other parties in an interaction), the  
23 “reputation” of a user can increase. (see the Kasbah system, MIT) [Kasbah 98] This  
24 reputation system is appropriate to a pseudonymous environment. Notice that gaining  
25 negative reputations is not useful when users can simply change identities. In one variation  
26 we can “block” certain users from providing information, when those users have negative  
27 reputations. Clearly, collaborative filtering or other data mining techniques could usefully  
28 allow for reputations when weighting information about an object.

### 29 30 **2.2.5 Meeting Planning** 31



A couple of concrete examples includes for example a real estate developer who develops clusters of home building sites or town homes which reflects the profiles of users sharing common or compatible profiles (e.g., socially, professionally or commercially) or recommending hotel lodging sites for individuals who share similar commonalties and happen to be visiting the same location or city. It is also possible to physically locate vendors at locations in which their most preferred customer prospects are most commonly physically traveling past or are physically situated (using LEIA) even considering where these prospects are during periods in which their shopper interest (or mood) for those items/services are heightened (or more generally factoring in both location and temporal (mood) factors into the user profile data as it is herein applied for this purpose. Whether users or vendors (subject to user consent) share customer data, there are other potentially interesting and relevant applications, for example, based upon LEIA a user's calendar schedule perhaps even verbal clues, it may be possible to coordinate meeting between users and professionals who can provide a useful service e.g., if a specialty physician happens to be situated in the user's present or anticipated locality and the physician would be available for an appointment at that time and location. It is possible that ad hoc use of shared clinic space may be available for such ad hoc appointments.

## **2.2.6 Investment Matchmaking/Venture Capital**

This section describes a market place in which start-ups can propose a venture or other financing needs in order to elicit an auction between different Venture Capital funds. We can allow potential investors to leverage securely confidential information regarding the details of business present and anticipated strategic alliances and customers. Each company securely registers all of this confidential business information SDI. It is even possible for confidential information about future R&D initiatives proprietary know how and intellectual property to be entrusted with SDI such that potential synergies may be determined well in advance of the market.

A domain expert(s) within SDI then determines potential synergy's between the various commercial entities both for purposes of facilitating introductions on a customer/vendor level, strategic partner level, as well as what particular financing sources which are compatible for the level of the financing needs and other characteristics of the financing

such as terms, involvement on a control and/or management level (as well as compatible commercial technology venue in which they participate). The key objective is to identify as first priority the commercial entities and customers which can provide the greatest degree of benefit to the vendor then to, secondly, determine which entities within the VC/financing community are able to bring the most valuable of these contacts to the table.

An investor may fund a small start-up and also improve the value of the start-up by facilitating a larger vendor becoming a customer. An investor might negotiate a special deal for such a vendor. In one variation SDI could be used to mediate the introduction of optimally strategic investment opportunities to an investor *and* to identify strategic synergies among and between different commercial entities who are already funded by an investor(s). In this way it may be possible to introduce these investors and create joint participation with the new prospective investor.

By leveraging SDI, the investors may even identify certain intriguing facts such as that the synergies between one of their own customers and the other investor are also significant and/or even dis-synergies exist between their own investors and themselves which do not exist with other investors (e.g. and investment in a competitor or customer of a competitor) that it would be mutually advantageous for the investors to entrust SDI with the job of making appropriate introductions for re-distributing equity ownership of the investors into those other companies which are more appropriate from a strategic standpoint.

The motivation behind agreeing to these introductions of course must be bona fide, serious and not initiated with the malicious intent to capture otherwise confidential information (particularly for those investors at higher level financing levels such as higher level VCs, investment bankers and more generally investors who tend to be more minimally involved in direct management or marketing initiatives for the company) it is certainly reasonable to imagine that some of these investors may wish to apply some of the above suggested techniques for finding (or even exchange equity opportunities within the commercial market partially for the purpose of diversifying their risk by thus owning smaller equity stakes in more companies and particularly those which are most commercially strategic and mutually synergistic in nature.

The powerful data collections of SDI regarding comprehensive commercial/investment data which of relevance to predicting the nature of an market dynamics of the associated

1 commercial entities, and to build insurance for an investor. The primary objective is to  
2 identify investments which are predicted (using a variety of methods both human intuition  
3 and empirically based) to behave very similarly (including e.g. possessing a very similar,  
4 upside and risk potential). A secondary objective (unlike the above variation) is to then  
5 evaluate and actually identify competing companies to the one(s) the investment in which is  
6 being insured where it is believed that the success of the competitor(s) would have a  
7 negative impact on that of the company. Thus risk is effectively diversified among the  
8 overall success of those companies. The insurance could include a deductible and premiums  
9 which are determined in accordance with the amount of coverage which the investor wishes  
10 to acquire. These premiums could be either paid in the form of cash or equity. In one  
11 variation "success" for insurance purposed could be measured as the success relative to the  
12 average of all of the others within the same similar group.

### 13 **2.2.7 Ideas Market**

14 Individuals can submit an idea to SDI along with the application framework, target market,  
15 or a list of vendors, which may most benefit as a result. Interested parties can then *bid* within  
16 SDI for the ideas. The system of SDI can *operate the auction privately*, and only provide  
17 bidders with a limited amount of information about the idea. The problem is to  
18 automatically appraise the value of an idea for a particular vendor without providing the  
19 vendor with the idea, so that when the auction is over only the agent that wins actually  
20 learns detailed information about the idea. The price paid may include an up-front price and  
21 promises of long-term royalties or equity, etc.

22 The privacy-secured ideas market is useful, because traditional negotiation processes fail:  
23 there is always a powerful motivation on the part of a commercial entity to steal the idea and  
24 leaving the originator with no compensation, consideration or acknowledgement. With SDI  
25 an individual with an idea is able to place an idea in the marketplace, so that commercial  
26 vendors can provide SDI with information regarding their business, to allow SDI to value  
27 the idea autonomously on behalf of the vendor; i.e. on the basis of current and future  
28 commercial research initiatives. SDI computes the estimated requirements and economic  
29 value that any given submitted idea would have towards that vendor.

30 Human experts may play a role in evaluating the value of an idea to a vendor. In a busy idea  
31 market there is the need for a scheme by which ideas are automatically routed to the most



appropriate individual experts to evaluate that idea, which is either within SDI or within the organization which most likely needs the idea. This routing scheme could be based partially upon attributes associated with experts and features extracted from the description of an idea.

### 2.2.8 Negotiation Intermediation

Given the above general application framework for SDI in which commercial entities can securely entrust to SDI with their commercial marketing and technology related focus and objectives, it is a reasonable extension to further extend this data Rich framework to enable the above described exchange of confidential information among and between commercial entities to occur not merely in a purely general and non-dynamic fashion but to additionally incorporate timely information which may represent time sensitive and critical decision processes which are in the process of occurring at that particular time and which may be relevant to the mutual strategic discussions and negotiations between the entities. A very important application/example of this approach is the fact that commercial entities may be in the process of negotiation with another commercial entity regarding particular commercial or technology related opportunity very often (the majority of instances) the nature of these discussions and particularly the identity (identities) is maintained as strictly confidential information throughout the discourse of these negotiations. Though the following solution may not be appropriate for all cases of confidential discussions and negotiations there may be particular instances in which a knowledgeable “expert” within SDI could be confidentially kept apprised of the details of the negotiations as well as progress and any obstacles as they occur. Of course, SDI and the expert have additional knowledge about the general and temporally specific needs and objectives of other third parties who may be potentially interested in the nature of the deal in progress.

At a general level there may be the opportunity for the expert to provide high level probing questions to the appropriate parties or decision makers within the other commercial entities which certainly do not reveal any identifying information about the negotiating parties, do not reveal any proprietary technical details which would compromise the proprietary nature of this information and enable the identity(identities)of the parties to be indirectly inferred in this regard. However, business terms and perceived value exchange may be thus conveyed

as limited facts in the abstract or possibly in more detail depending on the particular situation. There may also be particular critical junctures and impasses in the negotiation at which point the vendor may be compelled to concede. These points may be critical points at which it may be opportune for SDI to initiate or go into a deeper level of investigation with the other parties regarding the critical issues at hand. The above precedent may even be replicated among all other third parties, which SDI deems to be potentially receptive and legible to the prospective deal. This process is analogous to time of purchase competition in that if/when another potential deal is identified from another entity which is more desirable the original negotiating party may be informed (as well as possibly the other relevant parties). This process may be reiterated as well.

### 2.3 A Real Time Experts Market

In this section we suggest a market for expert opinions, in which users with an information-need are matched in real-time with “experts”, that are prepared to provide information and opinions in return for a payment. This economic approach can be used for example, if a vendor would like a user to provide feedback about its products and/or services: the user becomes an “expert” and can receive payment in terms of discounts in the future. The information in SDI can be used for the automated selection and user targeting of tasks, based upon profile information.

We can allow users to be identified by another individual in which there are mutually beneficial opportunities for both parties to interact and terms/conditions for the disclosure to the other party is defined within the user’s data disclosure policy. For example, issued patent 6,029,195 Herz, et al System for Customized Electronic Identification of Desirable Objects describes a “system methodology by which users are able to find a knowledge domain expert to answer a query, deliver personalized advice for a particular issue or problem to which they are extremely knowledgeable about, and provide references to other information sources.

The parent issued patent U.S Patent # 5754939 describes techniques for identifying experts on a communication network based upon their profiles and the search profile of a requester. In the present methodology we further provide an economic mechanism to encourage well qualified experts to provide options, and find experts. There is currently within the prior art economic incentive schemes which can be adapted and effectively

employed in this case in which it is possible to reward very well qualified (and truthful) opinions on a very broad range of issues and domains in which the accuracy of the ultimate outcome of a particular opinion is directly rewarded. This mechanism effectively insures the quality of the referrals/recommendations and is considerably more efficient than immediate payments for referrals where quality and truthfulness of the expert advice cannot be audited prior to payment, thus creating a disincentive to provide truly “valuable” advice.

The system of SDI can *forward requests for information to experts within the same system as time-of-purchase competition, SDI acts as a CLEARING HOUSE for requests for information, and experts can bid to provide responses*. The system of SDI can also help users to choose between experts, based on proficiency profiles that are derived from the value of information that experts have provided to other users. We can allow experts to bid for the right to annotate and provide ratings for particular pieces of information, products, services, etc.

In a web annotation example, as users provide annotations and recommend links, and other users provide feedback about annotations, then the system of SDI can build a “proficiency profile” for a user, to indicate the ability of a user to add value to the browsing experiences of other users. The information retrieval and document clustering methods as taught within the parent case #57549398 provide a statistically sound methodology to develop a user profile that predicts the “proficiency” of a user to provide recommendations about objects that fall in particular clusters.

In the economic variation we can allow agents to bid for the right to receive high quality recommendations, at least in the case where recommendations are new and cannot be duplication on the basis of current recommendations. We can create an “experts market”, where experts are assessed on their ability to provide quality to previous agents, as noted by those agents. SDI could also suggest a list of content categories that an expert might consider operating within. We can also apply the rapid profiling techniques in the parent case to assess a user’s expertise based on his/her response to a set of questions. A rapid profiling tree is essentially a decision tree, which can be used to present a sequence of items to which a user “responds”. The rapid profiling method enables a user profile to be constructed in the shortest possible sequence of inquiries to the user. In this case, the

response is the user's recommendation links (which are then judged). With the benefit of a comprehensive proficiency profile of the user the system may identify and present items to the user automatically (if s/he is amenable) for which the user is predicted to be proficient.

The market for experts and real-time information filtering can be applied to new content.

It can further be combined within the framework of the present incentive scheme to encourage users who tend to individually find news, which is of particular relevance to their daily lives and unique interests to proactively collect information (including but not limited to digital recorded audio/video) and provide that information to news distribution systems. Collaborative filtering leveraging both the overall quality/interest of user content as well as of the nature of the new content which that particular user is providing can provide a nice method to filter and channel this information.

The present clustering techniques may also be used to identify users that are able to provide useful new content. This content may range from commentary, opinions, critiques and





1 The criteria for matching the seller with a qualified prospective buyer is represented by  
2 client-provided meta-information associated with profile information, and can be used  
3 within a collaborative filtering system to determine the level of predicted interest which the  
4 buyer is likely to have for the items offered by a particular seller. The product space  
5 predicted to be useful to a buyer can be predicted based on his/her profile, browsing  
6 patterns, etc. Context can also play an important role, for example is the buyer currently  
7 engaged in browsing or buying related activities, can the buyer be persuaded to buy now  
8 with a good offer, or will the buyer never buy until he/she has performed more product  
9 search?

10 Vendors may also compete for sellers based on the profile information and success-profiles  
11 of a seller, and a perceived match with the products offered by the vendor. A seller may also  
12 work for numerous categories of products/services and vendors. Finally, the sellers in the  
13 market could enter into a market-based system, so that buyers bid for the right to work with  
14 a particular sales person.

15 Two other aspects of the present access to or by a prospective sales person include

16 (1) In an off-line context, we can use a LEIA-based method to identify buyers and sellers  
17 with similar profiles in physical space, and dynamically reroute their paths to allow a  
18 meeting to discuss a possible trade.

19 (2) Personal "Chemistry". In addition to a pure "product-space" set of profile features, the  
20 system might also consider wider compatibility between sellers and buyers in making  
21 meetings, for example choosing to introduce agents that share similar hobbies and spare-  
22 time activities.

23 The problem might also be informational: e.g. find an expert on ancient American  
24 civilization for purposes of writing an article, or answering a specific question. Relevant  
25 information might include the expert's resume, and the expert's knowledge expertise profile  
26 developed from his/her activities in responding to previous queries.

27 We might use a "fuzzy rule" to determine whether a user has a profile that is sufficiently  
28 close to an agent's goal profile to allow an agent-agent interaction and exchange of  
29 information. When the rule does not quite fire the system of SDI might also seek to clarify  
30 points of uncertainty, requesting further information until there is enough information to  
31 decide on the appropriateness of a contact.

1 It is clearly desirable to automate information exchange as much as possible, so that the first  
2 thing that a user knows about an agent-agent negotiation is after a deal has been struck that  
3 satisfies a user's preferences.

4 One approach is to perform data analysis on a large data set of users who share very similar  
5 profiles and to perform very similar actions and behaviors in all aspects of their successful  
6 interactions with other agents. The data analysis might also be used to suggest to agents  
7 when there is a problem, and when it will be useful for a user to provide more accurate  
8 information about what types of actions it is looking to perform in certain scenarios, i.e.  
9 when the current level of confidence about the action desired by the user is low. In many  
10 cases the system might be able to identify with a high statistical confidence level an  
11 appropriate action based on the collective statistics of other similar users and conditions.

12 When uncertainty exists the system can suggest the actions that seem reasonable, and allow  
13 the user to choose one action. A natural language interface might communicate to the user  
14 the system's interpretation of the user's current behavior, the statistical correlations with  
15 recommended autonomous actions.

16 For example, if a user is at home in his/her study between 9-11PM she/he is likely to be  
17 engaged in work relative to a research publication thus an invitation to other colleagues in  
18 this area will be extended and/or communication requests accepted to contact the user  
19 accordingly under these conditions. The system may detect other conditions in which the  
20 user may reject requests to correspond with these individuals, when the user's current  
21 context is inappropriate. Given this kind of feedback about actions, a user may now  
22 periodically review and adopt rules. For example, a user could also state that whenever a  
23 user has recently corresponded with a particular individual then the user is in a social mind-  
24 set, and would rather not engage in work-related activities.

25 There are numerous potential inputs to the system which could be considered in predicting  
26 what agent-mediated actions are, in fact, appropriate. The techniques of the parent patent  
27 application [INSERT US PATENT NO] allow us to make strong inferences about the  
28 particular mind set of a user; i.e. reflecting present interests or preferences which the user is  
29 likely to be receptive to, presently based upon such clues as who the user is presently  
30 interacting with, the content profiles of the present real-time dialogues, e. g. typed or spoken  
31 through a communications network, (or simply passively collected off-line), the object





currently viewing, thus enabling him/her to more efficiently and intelligently assess the needs of the customer.

This methodology could be further enhanced through the tracking of the user's pseudonym and associated profile data provided via SDI, which represents a far richer set of information than that used by eStara (which only makes use of the current web page identifier). A customer's full profile can include demographics, previous purchases, previous web-site visits, physical measurements (for clothing purchases), health history, and income.

In a sense, individual sales representatives also have personal profiles; these can include experience, demographics, languages spoken, previous customer contacts, and product knowledge. An especially important factor is a representative's relative success in selling particular products – this history may reflect various personal strengths and weaknesses.

When the customer clicks on an on-line catalog's "talk" button, his profile is immediately forwarded to the sales system, along with the particular product he is viewing. Matching algorithms (as described elsewhere in the patent) then choose the most relevant sales representative currently available for on-line communications. The representative is presented with a screen containing facts pertinent to the sale (including product details and extracts of the customer's profile), and he is then linked to the customer via Internet telephony.

Extensions to this technology include:

### 1. Customized Price Discounts and Promotional Offer Recommendations

By combining the predicted affinity of a user toward a particular item with price elasticity metrics (gleaned from his/her transaction logs and from real-time feedback provided to the system by the sales representative (e.g., the customer appears to be in a

happy and spending mood)), it is possible to calculate a “personalized” price that optimizes the expected profit from a sale. This optimized sales price is communicated to the sales representative, who is authorized to offer the discount to the customer. Digital coupon technology would be used to ensure that the offer is redeemed by that customer exclusively, and within a pre-set time limit.

## 2. Automated Sales Representative Support

While interacting with the customer, the sales representative is presented with constantly-updated screen that provides information to support his sales effort.

### a. Rapid Profiling

If not much information is known about a customer, the representative would be guided through a script designed to elicit the maximum amount of information in the least number of questions. This information would be used to rapidly supplement the customer’s profile.

### b. Customized Sales Scripts

The representative’s scripted sales pitch would be dynamically adjusted to reflect the profile and current attitude of the customer. For example, if a customer is having trouble making a decision to buy (which could be detected by voice analysis tools and the total time elapsed in the transaction), especially persuasive text, geared to the customer’s demographic, would be presented to the representative.

### c. Additional Offers

If separate, but related, products are predicted to be of special interest to the customer, the sales representative is alerted so that he can promote the additional offers while still in contact with the customer.

In this variation, SDI can act as an exchange where professionals can exchange personal favors, for example one agent might indicate that it wishes to be introduced to another agent that can help with a particular professional problem. Both buyers and sellers can submit their respective needs, for example the need for a particular personal or professional favor, and in the case of sellers the ability to provide favors. Other personal information, for example: professional, political contacts, organizational affiliations, areas of competence, professional responsibility and spheres of influence etc. can be confidentially entrusted with SDI and used to make particular types of matches.

SDI can initially automatically define appropriate matches between the “buyers” and “sellers”. Given sufficient data, humans may use content analysis techniques to define rules based upon certain key examples or commonly occurring request and matching solution pairs. The techniques of the parent patent application US Patent # 5754938 can also be used to refine rules, and create new rules, for the cases in which the level of confidence with introductions based with the current rule set is too low. The mechanism for establishing the value at which favors are sold might involve a bidding scheme, but in this case professionals may bid a “personal favor” in return for another favor. I.e. the transaction becomes disintermediated from the outside economy, with a pure non-money based method for exchanging favors. A user with the ability to provide a favor receives bids from other users in need of the favor she/he is able to provide. Of course, the value of a given favor is likely to be worth more to some individuals than others.

The nature of the favor may range from very small and mundane, e.g., providing professional or personal advice or answers to a question, or it may be very significant and resource intensive from the standpoint of the provider, e.g., making a professional recommendation for a job position or political decision or providing recommendations/referrals for clients on a professional level or providing business recommendations to a partner or affiliate regarding a given individual or commercial entity with which to do business.

The techniques of peer-ratings within a reputation system can be introduced, and extended to cover a wide array of qualitative descriptive attributes to reveal the qualifications and ability of an individual, to allow some individuals in the bidding process to be excluded in

1 an initial "filtering process", because they fall below expected standards of service and  
2 performance. An auction based model may establish a price for a particular type of favor,  
3 and for a particular individual, and therefore users with high performance ratings could be  
4 expected to sell for higher prices than other agents.

## 5 6 BARTER EXCHANGE FOR FAVORS 7

8 The qualification threshold of the provider of the favor is more important in some  
9 applications than in others. For example, if the favor involves an interaction with a  
10 particular individual for organization (e.g. a political favor) in which the individual's  
11 anticipated ability and qualifications may be a significant factor in the likelihood to achieve  
12 the desired results. And thresholds may be set relating to quality or qualifications in their  
13 regard depending upon the relevance of this factor and/or the importance of the favor, which  
14 depends upon the type of favor requested. The bidding process is typically a reverse auction  
15 in which the request favor is passed around to multiple individuals (or entities) which SDI  
16 deems qualified for the particular favor requested (thus the lowest price which fits the  
17 qualifications criteria is typically the offer which is accepted. In other (perhaps many) of the  
18 barter exchange for favors, the qualifications of the requester may be an important factor for  
19 the requestee. These qualifications may include, not only professional abilities,  
20 proficiencies and credentials, but potentially much more subtle attributes relating to the  
21 individual's interpersonal, social or psychological profile and/or behavioral profile, for  
22 example, how SDI predicts that the individual receiving the favor will integrate and adapt  
23 within a particular professional context and/or interpersonal context which defines the  
24 situation and/or environment with which that individual must interact and/or perform, e.g.,  
25 as in a business deal, a professional opportunity, an educational opportunity, social  
26 opportunity, etc. In order for this pre-qualification of the requester to unilaterally work so  
27 as to assure the requestee with the level of confidence and trust in the requester which  
28 approaches that of a traditional favor (in which the requester and requestee are typically  
29 extremely knowledgeable and trusting in one another) requires a very efficient and perhaps  
30 fairly comprehensive reputation system revealing much about the above described personal  
31 and professional aspects of one of both parties. As suggested, in such a system and just as

important, is the trustworthiness of the individual providing his/her opinion about the credentialed profiles of the parties. Indeed a political or decision maker in a large organization would be extremely remiss (and perhaps even politically harmed) if an office or position were provided to another individual (requester) in which weaknesses in that individual's professional and/or personality later became evident to the detriment of the organization or political entity, as an example.

It is even possible, that in some cases, the need may exist for a pricing model in which the individual providing these credentials through the reputation systems is compensated for the task. If such an individual must be fairly knowledgeable and assure to properly judge professional or subtle personality components, this individual may develop an independent reputation for effectively and truthfully judging these certain characteristics on a category or domain basis. If the subset of “judges” is fully distributed and ubiquitous permeating a substantial portion of the population as a default, if the qualifications of the judge are unable to be validated in advance, it may be preferable to provide payment following a reasonable level competency and success on the part of the requestee and requester. Or, in another variation, users are required whenever appreciated (and in accordance with their own judging qualifications) to provide opinions on the individuals (or other entities), and they must comply up to a certain quota and with truthful and accurate opinions in order to maintain privileges to engage in the service (e.g., [www.favor.com](http://www.favor.com)) or other applications or services resulting from the reputation system. It is also, of course, critical to ensure the judge possesses with regards to the individual or entity with which he/she is affiliated. In this case, the assumption is that the barter exchange for favors could cover nearly any type of favor desired by a user and the reputation system is implemented for most individuals. Of course, the present system may also be extended within a business context in which the favors include those provided by in between businesses. In both the individual and business application context, it is necessary for the above described use of privacy policies to dictate what types of individuals or entities a user or organization to include, exclude or price discriminations against e.g. to various degrees).

1 The following section "Agent-Mediated Value Exchange in the Supply Chain" talks about  
2 an economic system methodology in which a "value chain" is established for which entities  
3 may compensate other entities for present or often future value, which is likely to be  
4 sustained as a result of certain actions performed on behalf of that entity. This value often  
5 conveyed in the form of "barterables" could also be provided in the form of "favors"  
6 between the commercial entities (or potentially in a variation, individuals).

### 7 8 **2.3.2 A Market for Referrals**

9  
10 Situations frequently arise in a variety of contexts of human interaction (whether social or  
11 professional) in which a user may wish to refer the user they are in contact with to  
12 another individual. Often this occurs in a professional services context which a user has a  
13 particular need/or other characteristics which make him/her an appropriate match for the  
14 services provided by the other party. Or in a business context, often a user will forward a  
15 business contact or associate to another colleague who is deemed more appropriate for  
16 the particular context and/or scope of business. Likewise, in a personal or social context  
17 users may sometimes meet two or more individuals which they observe or perceive share  
18 common interests, goals or beliefs or perhaps possess complementary capabilities,  
19 knowledge, or characteristics. In each of the above scenarios, virtual tags may provide  
20 substantial benefits. For example, the referring user could forward the relevant portion of  
21 the profile and identified need of the user to the referring party whose user agent may  
22 determine the acceptability of the request and/or the priority with which a communication  
23 or meeting could be scheduled (e.g., as could be automatically arranged by/between the  
24 two party' calendaring agents). If the referring party's agent is unable to make a decision  
25 or priority assessment for scheduling purposes) on behalf of the user, the agent could  
26 instead try to contact the individual him/herself for assistance (and statistical feedback to  
27 the system's data model). In order for these types of referrals to be performed efficiently,  
28 the area of expertise required can be specified, and provisions can be made about the type  
29 of referrals that a professional will accept.





1 enable free access by users on the Internet who may want to, and are able to, usefully  
2 contribute. The idea behind the present scheme is to leverage techniques of natural  
3 language processing and/or information retrieval in order to develop profiles of  
4 individuals based upon the value of their past contributions to documents; for example  
5 informational content, style of authoring, etc. In a Web-based application the content can  
6 be graphical designed content, perhaps even video segments. Within SDI it is possible,  
7 based on feedback from previous articles, to anticipate how valuable the information is  
8 likely to be for a particular readership.

9  
10 The problem might also be informational: e.g. find an expert on ancient American  
11 civilization for purposes of writing an article, or answering a specific question. Relevant  
12 information might include the expert's resume, and the expert's knowledge expertise profile  
13 developed from his/her activities in responding to previous queries. Level of expertise  
14 might also include the size of projects performed within a particular specialized area, and  
15 relevant education qualifications.

#### 16 **2.3.4 A System for Smart Consumer Research**

17  
18 There is a significant need within the field of consumer product creation and development to  
19 be able to more instantaneously, and on a larger scale, collect direct feedback from a large  
20 number of consumers in lieu of much more sparsely populated focus groups and "product  
21 development" experts. There is also a need to better characterize consumers whose feedback  
22 is utilized in any kind of market test analysis. Collaborative filtering can provide a very  
23 efficient solution to both of these problems by enabling focus of feedback from selectively  
24 chosen consumers whose attitudes and opinions (and even marketing ideas) are the most  
25 meaningful and representative of the large segments, comprising the majority of the  
26 consumer population for that product. We can further extend this methodology to determine  
27 and measure the "value" of certain individuals in the process of product assessment, i.e. to  
28 identify the "experts" in a particular area.

29 The information in SDI can be used to identify appropriate clusters of users for a particular  
30 product, so that a prototype or actual product can be presented to users, and feedback  
31 collected. The goal is to collect detailed feedback: across a variety of quality and consumer



1 customized offers via coupons, through the use of cryptographically secure digital  
2 coupons which can authenticate the user and be delivered to the user in a non-transferable  
3 and non reusable manner.

4  
5 Within the system of SDI we can also use secure user features in the SDI database to  
6 offer special offers to an individual. Of course, the vendor that requests coupon  
7 generation based on information about profiles can execute the request without knowing  
8 the profile or identity of the user.

9  
10 As explained in the top-level SDI description, within the query-execution module rules ,  
11 can be specified that determine whether a profile is suitable for an offer, and the type of  
12 offer to make; and the system of SDI can automatically furnish the provider of the profile  
13 information with a coupon (as long as that is compatible with that user's requests). The  
14 offer to a user can therefore be flexible, based upon a user's profile. The user may be  
15 given an encrypted code to present to a cashier when it makes a purchase, allowing offers  
16 to be made by vendors that are not on-line, to encourage a user to drive to their physical  
17 stores. More information can be available within SDI, for example to allow a user to  
18 receive information about quality ratings of vendors, for the products and services that  
19 have been provided to other users. In return, the vendor may request certain purchase  
20 pattern data from the user. In another variation, the user might enter certain identifying  
21 information such as the first several words of the offer, and then receive an email or fax  
22 of a barcode to take to the shop for scanning.

23  
24 This method of discounting can be extended to users with Personal Digital Assistants  
25 (e.g. palm computing devices) in a store, that can use a scanner to enter bar codes of  
26 products, and then enter into an agent-mediated negotiation to secure a good price for the  
27 product. The user may be requested to disclose certain personal data, if a "better" offer is  
28 to be provided. Again, even in the in-store application the vendor does not need to  
29 receive actual information associated with the user, the processing can be performed  
30 remotely either at the user's ISP proxy server or centrally on a SDI server, with the

1 vendor providing methods to adjust prices and offers based on a user's profile, but  
2 without seeing the actual profile.

3  
4 SDI may also implement time-of-purchase competition, soliciting and revealing to the  
5 user competitive offers from vendors in stores in a close physical proximity to the user's  
6 location. The user could also be prevented with other types of useful information,  
7 including: information for competitive products including (nearby) locations, price and  
8 functionality for eliciting time of purchase competition as well as previous customer  
9 complaints about each vendor.

10  
11 Finally, an alternative to dynamic negotiation with the vendor is that price labels in stores  
12 could encode a *strategy* that the store proposes to use to price-negotiation with a user, and  
13 just download this strategy to SDI and combine it with the user's profile to compute the  
14 final offer. The user receives a validated offer from the vendor, that certifies that the user  
15 with his/her current pseudonym is able to receive the price or discount. The offer can be  
16 encrypted to prevent fraud by the user.

17  
18 As suggested above, SDI acting on behalf of the user may selectively release only that  
19 information which based upon the presumed price/offer generation model used by the  
20 vendor will elicit a price decrease delivered to that user. Within this price negotiation  
21 process we can require that SDI releases just the right of information to optimize a user's  
22 price. This can be part of the understanding with a vendor (even though a vendor may  
23 commit to a strategy up-front). SDI cannot falsify information on behalf of a user, but is  
24 free to withhold information.

25  
26 It is certainly conceivable that such a technique could be deployed by vendors, to allow  
27 customized pricing for users as they shop in a store via information encoded with product  
28 identities, and negotiated via profiles stored in decentralized SDI nodes. A useful default  
29 to make the system work for non-SDI based users could be to allow a user to user  
30 averaged ratings on various criteria, annotations, etc.





Another variation is an auction scenario, where a seller brings an item to market, and buyers compete for the item. A professional, that makes a living by buying items and selling them on, might like an exclusive SDI-based auction site. Within the system of SDI we can support a *network of auction-sites*, that push agents around between sites and revenue-share, with source web-sites providing shares of revenue achieved in subsequent purchases on other sites. In another variation, the operator of a small site may be linked from a larger site, and provide a share of revenues to that larger site. SDI has a couple of important roles: monitoring purchases and ensuring that contracts are fulfilled, and also estimating the value of links if an up-front price is to be negotiated, based on information about the profiles of streams of customers to a particular site.

The information in SDI can be used to allow disintermediation of consumer-to-consumer markets, with agents able to reach ideal prospects based on personal information stored in the shared database, make an anonymous offer, and consummate the trade of possible. This is likely the way of the future in this consumer-to-consumer market. Just as a vendor in current e-commerce systems can store information about its own user-base, and use that information to send personalized offers to users, within the system of shared information in SDI all agents can share information with all other agents, and the playing field is leveled. Users can query the pseudonymous user profile data and (if permitted), gain direct personal access to appropriate target customer prospect.

Conversely, prospective customers who are interested in a particular item(s) may query and access the database of users who are knowledgeable regarding that particular product/service (which may possibly be under a pseudonym), and also accessible to the user, i.e. available at an appropriate time and in an appropriate location.

## 2.6 Transportation Example

Another application of user profile database, in the case that the database contains also real-time information, is in an application to transport scheduling problems. The traffic systems in many large systems are congested. Via the system of SDI it would be possible to allow agents to represent the wishes of their uses within an SDI-enabled mechanism that controlled access to certain lanes on highways. The goal might be to control the flow of traffic, such that users agree to pool vehicles based on similarities across planned trips, and

1 users also payments in return for the right to travel at certain times and in certain lanes. The  
2 goal is a more efficient system that maximizes the sum value to all agents.

3 This is an example of a more general application area, where agents can look for  
4 opportunities to *change their actions and behavior in return for higher system efficiency,*  
5 *and/or payments.* We describe below an application to user agents negotiating on behalf of  
6 individuals ("travelers") to optimize the efficiency of a shared transportation system. The  
7 fundamental idea is to embed multi-user awareness into the basic fabric of a transportation  
8 system, in which the system aims to maximize the efficiency of the system by allowing  
9 payments between individuals.

10 To accomplish this task, the system must be able to anticipate both the immediate direct  
11 effects, a particular accommodative action will have on another agent(s) (mobile entities) as  
12 well as the indirect (cumulative) effect on the other agents and the associated time delay  
13 impact on each agent and on the entire transportation system as a whole.

14  
15 The techniques of the LEIA can be efficiently employed in tracking the present and  
16 anticipated activities, location and movement patterns of individuals who are in the process  
17 of traveling to a destination via simple or multiple modes of transportation, e.g., any  
18 combination of the following, i.e., pedestrians, automobiles, taxi, train, and bus (including  
19 public transportation). LEIA can be applied in both scheduling the flow of traffic for  
20 transportation media which are not subject to fixed time schedules (cabs, autos, pedestrians)  
21 as well as to dynamically improve overall efficiency of the movement patterns in order to  
22 assure that users on a collective basis arrive at their desired destinations with minimal delay.

23 The system of SDI can implement a real-time market where agents make and receive  
24 payments in return for changed actions from other users. For example, agents with high  
25 priority may be prepared to pay more than other agents for the right to enter a highway at  
26 any particular moment in time.

27 The market pressures at any point in time are between the marginal cost to an individual for  
28 changing his/her plan, and the cumulative marginal benefit to other agents in the system.  
29 Within incentive-compatible mechanisms, such as the regular price-based market place (e.g.  
30 the Walrasian tatonnement model) with enough agents, then it is optimal for an agent to



1 reveal its true value for different outcomes to the marketplace, and the system of SDI in  
2 “clearing” the market can maximize system efficiency.

3 The central contribution of the SDI system is that it acts as a trusted auctioneer, receiving  
4 information from agents, and implementing an outcome based on that information without  
5 releasing that information to other agents. All that agents see are requested actions, and  
6 payments received or payments to make.

7 We could model a core sample of individuals on an active basis to determine implied  
8 valuation functions for different types of users, based on feedback provided by those users  
9 about how happy they are at any point in time. Within an incentive compatible system it is  
10 optimal for an agent to have a truthful and complete representation of its valuation function.

11 The problem is to discover that valuation function. Any help from SDI, via data mining  
12 techniques (i.e. similarities between my profile and the profile of another user that has rated  
13 his/her happiness) can be beneficial in reducing the costs of participation in such a scheme.

14 Then, before taking a journey, the user could provide some explicit information regarding  
15 the nature of the travel, for example the level of importance in minimizing travel time,  
16 whether the trip is business or sightseeing, the time of day, etc. With this information the  
17 system can assess the user’s valuation function, and then have the user participate within the  
18 market, and secure a plan based on actions from all agents. The system might compute price  
19 thresholds, that represent different things the user can achieve for different prices. The  
20 system based upon its determination of situations/context can presents the price and time  
21 and urgency inferred for a user prior; which can be reviewed and corrected by the user.

22 With this profile information, the auctioneer can now compute prices to clear the market and  
23 maximize economic efficiency in the system; taking the following approach:

- 24 1. Armed with statistical information about the users of the transportation system e.g.  
25 the highway for the next hour, compute static (fixed) prices for different options.  
26 The goal is to maximize the efficiency in the market through a simple pricing  
27 mechanism, that is computed based on information provided to the auctioneer in  
28 advance.
- 29 2. Announce the prices, and allow users to use the system as they desire, with  
30 appropriate payments made. Payments can be readily extracted from agents via  
31 automatic toll systems.

1 Of course over time the performance of such a system could be optimized, as the auctioneer  
2 (the SDI marketplace) learns about consumption patterns of the agents over time. One way  
3 to provide incentives for agents to provide information to the system to allow up-front  
4 computation of useful prices to support an efficient outcome would be provide *discounts* to  
5 the agents that provide information. When such a system works well then roadway traffic  
6 can be controlled, congestion controlled, via an automated price-based system. Additional  
7 information made available from within SDI might allow adjustments to prices, based on  
8 unexpected flows of traffic. Just because the prices are optimized once up-front, based on  
9 projections, and announced to agents using the system; it can still be possible to adjust  
10 prices during any particular period—so long as there is a method for that feedback to alter  
11 the actions of agents, i.e. there must always be an alternative to make such a system work.  
12 Consider a two-road system, one shorter and one longer, then the longer route would be  
13 priced less than the shorter route, and the price differential adjusted based on current levels  
14 of congestion.

15 Again, to clarify, SDI has a key role in this system. Agents are only prepared to reveal  
16 valuation and trip data within a system that carefully protects their privacy, and controls that  
17 information.

18 Via the optimization performed to compute appropriate prices the model takes into account  
19 the effect on congestion that consumption decisions have, i.e. there is a relationship between  
20 the value to an agent that selects option A and the number of agents that select option A.  
21 The auctioneer needs a model of the transportation system itself to perform this type of  
22 computation. Essentially, we compute the fixed point of a non-linear system of equations.  
23 Stochastic optimization techniques are suitable for such a calculation.

24 As the system collects more data it can be less important to require explicit information  
25 from agents within the system, unless it is believed that there is a special situation about to  
26 occur (for example a Flyers game), which will have unusual effects on traffic patterns.

27 Co-pending patent application entitled LEIA-TR describes flexible tools and automatic  
28 traffic ticketing which could be integrated in to the same transaction infrastructure. SDI  
29 would provide a methodology to elicit information from agents, and dynamically set the  
30 price of traffic tickets and parking meter tolls to maximize the performance of a City's  
31 parking resources. For example, one goal of the transportation authority can be to optimize

the amount of revenues it can achieve, by charging more during “high-demand” travel periods.

Extensions of the present scheme could include coalition discounts, where coalitions of users can negotiate discounts based on group purchasing power (let us all use your system for a 10% discount or we will use an alternative road system). The coalitions may be formed automatically within SDI using shared profile information, allowing agents with non-competing interests to “pool their buying power”. Finally, it is interesting to note that users who share identical travel objectives can be encouraged to share vehicles, and/or public transport can be provided as alternative means of transportation on a dynamic basis.

### 3. Client-Side Data Mining Applications

Finally, we describe applications of *client-side data mining*, where the private information about an agent remains on its local client machine. Personalized of services and products is performed by pushing methods to the client machine, where the methods compute based on local profile information the most appropriate information for the user. This is very useful because the user retains absolute control over his/her profile information, but can still receive the benefits of personalization. This is an extension of iamworthit: users still provide profile information to the central database, to allow central queries and to receive value for that information, but never release identified information to a vendor.

#### 3.1 Client-Side Personalization

Client-side personalization allows vendors to push personalization rules to client machines that are special queries which use information store in a local client profile database to provide a user with personalized information and/or services. The personalization rules use locally stored profile information to compute an optimal product or service, or to allow a vendor to configure a virtual shop front. The answer to a query is returned to a vendor, to allow the vendor to push suitable commands to the client to enable appropriate displays to be created on the device with which the user interacts with the client.

Client side personalization is useful within the system of SDI because:

- a) It allows vendors to leverage profile information across multiple profiles that belong to an individual without explicitly receiving the information.
- b) It provides users with the convenience of using their true identities in payment and specification of delivery addresses for goods, because vendors do not need to





16 One-time Client-side personalization. In this variation the agent computer only  
ever performs a single user-personalization step, which is then maintained by the  
user as part of its profile for that vendor and passed in the future to allow  
personalization.

It can be computationally complex to compute the result of a personalization query. For example, profiles about the objects on a vendors web page or in its inventory can be as large as the object descriptions themselves, and full decision tree representations to decide how to target a consumer can be very large and complex. In cases where this is a problem, and it is inefficient or infeasible for the vendor to push a complete decision tree to the user we propose the following solution:

17 Iterative Client-side Personalization. In iterative client-side personalization the  
18 vendor and the user, via their agent computers, participate in an iterative  
19 distributed protocol to compute a personalization result for a user based on its  
20 profile. A similar method was earlier disclosed in US Patent #5753938 The idea  
21 is to structure the decision tree, and for example pass initially the first few levels  
22 to the agent. The agent computes the result of those levels, responds to the  
23 vendor, and the vendor passes the next few levels that are relevant to the result  
24 received. In this way, only the parts of the decision tree which are actually used to  
25 compute the result of a personalization query are exchanged with the consumer's  
26 agent computer.  
27

Specifically for collaborative filtering applications, a simple technique can be used to reduce the amount of information that must be provided to a user to compute the result of a personalization query. In collaborative filtering a vendor's decision about which objects are

1 appropriate to show to a user is made on the basis of a similarity comparison between the  
2 user's profile and the profile of other users for which the vendor has information about how  
3 they responded to certain objects. The straightforward method is to pass the profiles for an  
4 entire user population to the agent's client machine. This is undesirable when there are many  
5 users in the population.

6 A better solution is to pass only cluster centers to the client, where the cluster centers are  
7 chosen to allow collaborative filtering to a sufficient degree of accuracy. For example, we  
8 illustrate below an efficient method for collaborative filtering at the client, and our proposed  
9 efficient implementation. This is illustrated in Figure 20.

10  
11 The idea is to select a subset of users that represent the entire space quite well, which the  
12 vendor has collected by performing data mining queries on the central SDI data warehouse,  
13 and then submit just those data points to the client. In the Figure we plot the complete set of  
14 user profiles in profile space, with each user profile associated with information about how  
15 to target a product or service to a user with that profile. A typical method to perform  
16 collaborative filtering with data of this kind is to find profiles in the neighborhood of a new  
17 data point, illustrated with \* in the above diagram, and then base the decision on what was  
18 successful or unsuccessful for those users. Clustering replaces a set of similar users with one  
19 "cluster point" that represents what information a typical user in that region should be  
20 shown. This can be computed using standard cluster analysis techniques. The reduced space  
21 of data points, shown on the right, is sent to the client, and the client performs collaborative  
22 filtering over the cluster points to computer an approximate solution to the full  
23 personalization query. Completing the example, the agent computer can determine which  
24 "aggregate user" its profile is closest to, for example Mary or Yu-wen, and send this to the  
25 vendor.

26 Notice that the vendor does not even have to provide the "y-values", or the personalization  
27 methods which are associated with each user profile, it is sufficient for the agent to compute  
28 the closest cluster center and pass that information back to the vendor. This is useful to  
29 vendors because the information is not useful to anyone but themselves, because no-one else  
30 can interpret what it means to be like aggregate "Mary".





Clearly, standardization of user's profile information is necessary, so that all vendors can write methods that can access a user's profile. One solution that would support implementations from multiple providers of profiling functionality (i.e. would not require clients running standard SDI-profiling software) would be to provide "XML-data wrappers", that convert non-standard formats into a standard XML data format, that is used by all personalization methods. One language for representing user profiles has been proposed by the W3C consortium, known as PIDL (personalization definition language) [W3C-PIDL].

### 3.2 Real-Time Targeted Advertising

A useful application of SDI allows vendors to provide targeted advertisements to users based on their browsing behavior and other profile information, but without actually receiving explicit information about a user’s profile. We propose an advertising-server, which controls the network of adverts on behalf of vendors and users. A user’s client machine can receive a choice of adverts whenever it hits a web site in the ianworthit advertising network, with local evaluation based on the local profile information at a client to decide on an appropriate advert to display.

The system is designed to support dynamic personalized advertising, with vendor-side competition for the right to present an advert to a buyer, and buyer side criteria for accepting or refusing an advert. A technical solution is presented to allow vendors to select an appropriate bidding-policy, to maximize revenue from advertising. We use client-side profiling and advertising selection to support targeted advertising with absolute guarantees



1 have a more complete picture of a user's profile, and can compute the expected value of  
 2 viewing an advert based on the advert and the price offered by vendors. In general, as the  
 3 acceptability of an advert increases vendors will wish to bid more, and users will accept the  
 4 advert for less money. Competition between vendors drives the bid price up, and allows  
 5 users to receive the value of information about their profile to advertisers.

### 6 **3.2.1 Example: Technique to compute the expected value of an** 7 **advert.**

8  
 9 A collaborative filtering system can be used to predict the probability that a user with profile  
 10  $x$  will respond to an advert. The computation is based on historical information in the central  
 11 SDI data warehouse for similar users to the advert, and also on the basis of historical  
 12 information for similar users to similar advertisements if there is little information available  
 13 about the actual advertisement. Information can be provided to advertisers anonymously in  
 14 performing queries, and also randomized if necessary.

15 The basic technique is to select a set of users that have been shown the advert, or a similar  
 16 advert, and are close in profile attributes to the current buyer. Call this the "relevant set".  
 17 Given this, the probability  $\text{Hit}(x, \text{Ad})$  that a user with profile  $x$  will hit an advert with profile  
 18  $\text{Ad}$  can be computed as the ratio  $\text{Hit}(x) = \text{Num\_Hit} / \text{Num}$ , where  $\text{Num\_Hit}$  is the number  
 19 of users in the "relevant set" that responded to the advert, and  $\text{Num}$  is the total number of  
 20 users in the relevant set.

21 Although subject to a certain level of necessary uncertainty, because buyer behavior cannot  
 22 be predicted perfectly, ultimately it is useful to predict with as high an accuracy as possible  
 23 whether a buyer will respond. Given a probability  $\text{Hit}(x, \text{Ad})$  that an individual buyer will  
 24 respond to an advert, a vendor can define a bidding policy. The policy maps  $\text{Hit}(x, \text{Ad})$  to a  
 25 value to bid for the right to target a buyer.

26 Assume in this section that the profiling system places buyers into one of several "classes"  
 27 of buyers,  $C_1$  to  $C_n$ , given an advert with profile  $\text{Ad}$ , where each class has an associated hit  
 28 rate, i.e.  $x \in C_1 \Rightarrow \text{Hit}(x, \text{Ad}) = \text{Hit}(C_1, \text{Ad})$ . Furthermore, assume that the system also  
 29 predicts the average amount spent by a user that hits on the advert, or the average value to  
 30 the vendor from a hit, this can again depend on the class of buyer, and can be determined  
 31 within a profiling system based on historical information about this advert, or about similar

1 adverts. Let  $\text{Rev}(C1, \text{Ad})$  denote the value to a vendor with advert  $\text{Ad}$  of getting a hit with  
2 a buyer in class  $C1$ . Now, the vendor can determine its expected value for placing an advert  
3 to a user in each class, for example using the computation  $\text{Val}(C1, \text{Ad}) = \text{Hit}(C1, \text{Ad}) * \text{Rev}(C1, \text{Ad})$ .

5 The system of vendor-side advert competition works as follows. Once a user hits a web page  
6 with an ianworthit-targeted advert, a choice of adverts are pushed to the user's client from  
7 the ad server. Each advert is represented as a three-tuple:

8 ( link to location of advert, bidding function  $\text{Bid}(\text{hit}, \text{rev})$  , profiling function  
9  $\text{Profile}(x, \text{Ad})$ , value function  $\text{Rev}(x, \text{Ad})$  )

11 The profiling function  $\text{Hit}(x) = \text{Profile}(x, \text{Ad})$  is used to place a buyer into the appropriate  
12 class of buyers, and compute the hit probability for a buyer,  $\text{Hit}(x) = \text{Hit}(x, \text{Ad}) = \text{Hit}(C, \text{Ad})$ ,  
13 where  $x$  is the profile for the buyer,  $C$  is the class that the buyer's profile places it  
14 within, and  $\text{Hit}(C, \text{Ad})$  is the probability that a buyer in the class will respond to the advert.  
15 A possible implementation of the profiling function is to encode it using prototypical cluster  
16 centers for a buyer in each of the set of buyer class types, and then assign the buyer with  
17 profile  $x$  to the cluster that is closest (in some well-defined metric) to the cluster type. The  
18 metric does not need to be linear in each of the dimensions of a user profile, and in particular  
19 would be expected to be place more weight on terms that are important to the hit rate in a  
20 particular class of buyers.

21 As described, this method has the following useful characteristics:

- 22 a) the hit probability for a buyer is computed at the client machine, using the profile  
23 that is stored at the client for the buyer. The client machine does not need to  
24 release the user's profile, and the **user's privacy is assured**.
- 25 b) The information necessary to compute the hit probability can be encoded in space  
26 **linear in the number of clusters**, which is much more efficient that passing  
27 information about every relevant user profile to the client machine.

29 Without computational/communication restrictions one might pass historical information in  
30 the form of (Profile, Hit/Miss) pairs for users that have been shown the advert, or shown a  
31 similar advert. A nearest-neighbor algorithm could then be used to find the relevant set of

1 profiles for a new profile **x** and compute the expected **hit-rate** from the ratio of users with  
2 similar profiles that historically hit/missed the advert.

3 The cluster centers approximate this solution, representing the average hit-rate of “close”  
4 profiles so that the client machine can simply locate the closest cluster center and use that as  
5 a proxy for the probability that its user will hit the advert.

6 A similar technique can be used to compute the expected revenue from a buyer with profile  
7 **x** that hits an advert. In fact, this information can be computed using the same method,  
8 simply by associating an expected revenue with each cluster center.

### 9 **3.2.2 Client-side Advertising Reverse Auction**

10  
11 The client implements an auction for the right to show an advert to a user. The auction is a  
12 Vickrey auction between all the adverts that are passed by the ad server to the client. The  
13 Vickrey auction (Vickrey, 1961) is a second-price sealed-bid auction. Given bids **b1, b2, ...**  
14 **, bn** the auction sells the right to show an advert to a buyer to the vendor that bids the  
15 highest value for the value of the second-highest bid. The Vickrey auction is useful because  
16 it is truth-revealing. The optimal strategy of each vendor is to bid its true value for showing  
17 an advert to a user. For example, with a profit-margin of 20%, it is optimal for a vendor to  
18 submit **bid = 0.2 \* hit \* rev**. If accepted, it will pay at most **bid**, and it will pay only  
19 enough to shave the bid of its nearest competitor. The auction is sealed, so competitors do  
20 not ever see the information in failed bids. Cryptographic techniques can also be applied to  
21 ensure that the bids are not inflated by the auctioneer.

22 A user can also define an acceptance function in the auction, which represents her  
23 reservation price to view a particular advert. Assume that the reservation price depends on  
24 the hit rate **hit** of the advert, and represent the price as **Accept(hit)**.

25 The client-side advert auction runs as follows:

- 26 1. For every advert **Ad**, Compute the hit rate of advert **Ad**, given the profiling  
27 function **Profile(x, Ad)** and the buyer's profile **x**, and the expected revenue from  
28 showing the advert to the user, **rev**.
- 29 2. For every advert **Ad**, with hit rate **hit = Hit(x, Ad)**, compute the acceptance level  
30 **accept = Accept(Hit)** for the buyer and the bid for the vendor with the advert, **bid**

- 1       = Bid(**hit**, **rev**). If **accept** > **bid** then reject the bid for advert. If no adverts  
2       remain, jump to step 7 (in this case no advert is shown to the user).  
3       3. Given the bids that are not rejected by the client, sort them in order of increasing  
4       value, and insert a “buy” bid from the buyer equal to the value of its acceptance  
5       level for the advert with the bid of highest value.  
6       4. Accept the advert with the greatest bid, and charge the vendor the price of the  
7       second-highest bid.  
8       5. Fetch the graphics and URL for the advert, and display the advert to the user.  
9       6. The client monitors the actions of the buyer, and records (to be later transmitted to  
10      the iamworthit server) whether the buyer responds to the advert.

11  
12      The buy bid in Step 3 will be less than the bid for that advert, because the advert was not  
13      rejected in step 2. However, this buy bid might be greater than the value of the second-  
14      highest bid for an advert, and is required to make sure that the price paid by the vendor that  
15      wins the auction is greater than the buyer’s acceptance level. We make sure that the user  
16      cannot cheat by bidding just below the highest bid received by requiring that the user states  
17      his/her reservation value before the value of bids are revealed.

18      A more general system for dynamic customized advertising might allow a vendor to specify  
19      a bid for each type of buyer, and also a maximum budget, so that the vendor that places  
20      adverts can maintain control over its spending.

21      We might also allow users to specify in their advert acceptance policy how many adverts  
22      they are prepared to receive a day, so that they are not inundated with too many adverts,  
23      even if they receive financial compensation.

### 24      **3.2.3 Numerical Example**

25      The advert auction message from the ad server arrives with a choice of three adverts, **Ad1**,  
26      **Ad2**, and **Ad3**. Each advert is associated with a profiling function, a bidding function, and a  
27      value function. The client machine computes the hit rate for each advert, based on the  
28      profiling functions and its local profile for the buyer. Suppose **hit1** = 30%, **hit2** = 5% and  
29      **hit3** = 15%. The client machine also computes the expected revenue if the buyer hits an  
30      advert, **rev1** = \$2, **rev2** = \$8, **rev3** = \$3. The hit rate and revenue are used to compute bids

for each advert, using the bidding function. Suppose that the bids are **bid1** = \$0.50, **bid2** = \$0.70 and **bid3** = \$0.40.

Now, the client also computes the acceptance level for each advert, based on the hit rates predicted within the iamworthit system. Suppose **accept1** = \$0.30, **accept2** = \$1.00 and **accept3** = \$0.30. The bid for advert 2 is rejected because it is below the accept value. The bids for adverts 1 and 3 are accepted. Now, the auction is constructed with bids **bid1** and **bid3**, and **accept1** because that is the accept-value for the bid with the highest value that is not rejected. The auction takes the bids (0.50, 0.40, 0.30), and sells the right to advertise to the user to vendor 1 for \$0.40 (the value of the bid from vendor 3).

Finally, the client fetches the graphics and URL link information for the advert from vendor 1, and transfers payment from the vendor to the user for the right to show the advert.

### 3.2.4 Tuning an Advertising Strategy

Although the optimal strategy for a bidding agent in a single Vickrey auction is to reveal its true value for the right to show an advert to the user, the vendors are bidding over a number of different auctions, and might have a limited advertising budget.

#### *Selecting appropriate users.*

One good approach is to start with a reasonable policy and then adjust it dynamically, based on feedback received from adverts using techniques from reinforcement learning (see [RN 97] for an introduction). Off line simulation with methods such as Monte Carlo simulations can provide good initial policies, based on information about a simulated population of users, which could be provided anonymously by the system of Secure Data Interchange. Users can be provided with incentives to reveal information anonymously about their acceptance functions, to allow this type of off line modeling.

A vendor can tune an advertising policy with a random sample of users, simulating the auction that runs on a user's client machine. Metrics such as the average number of times that an advert is shown to a user of each type, and the average amount that the vendor pays to show the advert provide information to allow optimization. For example, the analysis might show that although a particular class of buyers are most likely to hit an advert, there is also a lot of competition to show adverts to buyers in that class, and the average cost to show an advert is high. In this case a vendor can conclude that it is more cost-effective to

1 advertise to buyers in cheaper but less relevant classes. The analysis can also be used to  
2 check that the average buyer is not setting an acceptance level higher than the vendor's bid,  
3 which again would indicate that the policy should target a different set of buyers.

#### 4 *Selecting appropriate adverts.*

5 The system that we have described can be extended to allow a vendor to select an advert to  
6 display to a user based on the type of profile of a user. For example, a vendor can pay for the  
7 right to be one of N vendors that compete in an auction for the right to advertise to users that  
8 hit the web page of a particular vendor. We can allow a vendor to not only tune its bid to the  
9 profile of a particular user that hits a web page, so that only users with a good fit with the  
10 service offered receive the web page, but we can also allow a vendor to tune the advert that  
11 it shows to a user.

12 We allow a vendor to submit a number of different adverts and bid functions via the  
13 advertising network server. The client-side auction can be expanded to allow multiple bids  
14 from each vendor, where each vendor can submit a number of bids, but only pays the second  
15 highest price bid from another vendor. Again, the auction is truth-revealing for a vendor.

### 16 **3.2.5 Automatic Advert-Replacement Systems**

17  
18 Although there are potential copyright violations in any system which alters the content of  
19 information published by a vendor before displaying that information to a user, we comment  
20 that the above described system can be applied even with vendors web pages that do not  
21 subscribe to the advertising network in SDI. Adverts can either be replaced, or added in  
22 separate windows. This enables any user that subscribes to SDI to receive personalized  
23 adverts and revenue from advertisers, irrespective of whether or not pages are within the  
24 SDI network.

25 The key problem in ad replacement is to identify an advertisement within a page, i.e.  
26 information that is not related to the core purpose of the page. At present adverts may take  
27 one of two forms:

- 28 (a) They can be statically, or dynamically, generated at the web server of the web  
29 page that the user hits, and pushed to the user directly at the same time as the  
30 general content.







time is achieved, across all segments of the highway system, as a result of each individual system recommendation to each driver. For example, an anticipated compounding congestion problem would warrant the vast majority of vehicles to take alternate routes until the congestion or predicted congestion is alleviated in that segment.

## 2. Use of LEIA for Creating Traffic Models for Purposes of Mapping Commercial Industrial and Residential Real Estate Market Opportunities –

A very useful application of LEIA involves the collection of traffic pattern data on a time-specific basis (the day of the week and time of day) for the traffic, as it passes each piece of real estate. An electronic map which is ideally Web based (potentially nationwide is generated and constantly updated based upon this data. Additional information may be provided which may include (but is not limited to):

a. Origin and destination information of the vehicular traffic (as captured by LEIA) which may, especially if correlated with time, suggest the nature and context of the driver's activities, e.g., rush hour traffic, errand traffic, etc. It may be useful to factor in the type of neighborhood the vehicle returns to every night, the type of commercial or business entity she/he drives to work to each day, etc.

b. Other activity-related clues which the user is willing to release, e.g., devices interacted with, content interacted with or transmitted information, etc., which may provide insights into the mind-set of which users tend to experience when in the vicinity of the real estate property.

c. User Profile Data – Aggregate purchase and content affinities as well as price elasticity data (gleaned from purchase statistics) could be very useful information to commercial real estate developers and purveyors. Users with the right profile and a receptive mindset are of particular interest.

1 It should be noted that the present system may be extended to residential real estate. E.g.,  
2 what types of jobs (such as quality of jobs) do local commuters have? What are their  
3 numbers? How far do they commute (particularly if they tend to commute further than  
4 the present real estate site)? Do their commuting routes tend to pass the current potential  
5 real estate site? The last three questions would also be particularly relevant as well for a  
6 prospective industrial real estate development opportunity.

7  
8 Industrial real estate developers also may be interested in mapped models of real estate  
9 depicting the professional and known likely educational characteristics of the associated  
10 local residents in that region? What are the other businesses at which they work? (If  
11 available) what are their particular positions/responsibilities?

12  
13 3. Use of LEIA for Creating Traffic and User Profile Models of Traffic Passing  
14 Billboard Sites and Providing a Map of Such Information on an Available Billboard Site  
15 Basis –

16  
17 The presently described techniques for providing dynamically updated informational  
18 maps containing detailed statistical data regarding vehicular traffic passing real estate  
19 sites can be further extended to similar maps of interest to advertisers which contain  
20 locations for available highway billboards. The present system further provides Web-  
21 based access, which enables advertisers to make reservations and purchases of such  
22 billboards. In one preferred variation, an economic model is deployed to optimally price  
23 the billboards. I.e., a varied representative sample of each type of billboard sharing  
24 similar traffic/user profile features with others is auctioned for this purpose. In another  
25 variation this on-line auction model is deployed for all billboards available by the system  
26 in order to provide a novel service to advertisers which is a “billboard auction” site.

27  
28 4. Use of LEIA for Enabling Drivers to Identify and View Sales Opportunities and  
29 Offers Associated with Physical Objects in the Surrounding Environment Utilizing  
30 Heads-Up Display Technology

The present system is an extension of the virtual tag methodology in which physical objects contain meta-data and in which the location of such objects with respect to the user is determinable either by transmission of a signal, which contains the exact physical coordinates of the object or a scheme by which this information may be pre-loaded, which is typically from a server, which contains such information (thus “non-wired” objects may be identified accordingly) to the user’s device, e.g., as she/he physically travels in the vicinity of such items. In either case, the present physical location of the object must be known in addition to that of the user (in the case of fixed objects, of course, the dynamic location tracking techniques are unnecessary). A primary commercial application of the present scheme applies to purchasable items. The user may submit his/her interest profile to SDI for purposes of being dynamically notified and disclosed of purchasables which match his/her user profile or specifically requested search criteria. The metadata may include a brief description and any additional level of information which the seller may be willing to disclose to that particular user. The seller may possibly request details regarding price elasticity of the buyer (purchase appetite) before quitting any discounts to the existing price. Prime example applications of the present scheme may include:

- 1...Revealing used car sales opportunities to drivers (by actually flagging relevant vehicles which the driver passes (or which pass the driver);
2. Real estate and home sales opportunities (as well as apartment and commercial real estate rental/leasing opportunities;
3. User-user introductions in which one or both users may be extremely interested in the other based upon matching criteria within their associated social, information and/or professional desires and assets.

The preferred underlying technology for the present heads-up display system involves a technique for tracking the orientation of the user's head as well as the direction of his/her pupils (it may be possible, however, to implement the system by detecting pupillary

direction based solely upon the position of the pupils relative to one another, using a 3-D optical tracking device. The iris of the user may also be utilized as a bio-metric identity of the particular user (see issued patents assigned to IrisScan Corp.) In the preferred implementation, the head-up projection device projects the virtual flagger upon the portion of the windshield or window(s) which is in direct line of site between the user's pupils and the physical object of interest. The technique may utilize heads-up windshield display technology and/or remote projection (e.g., for projection upon the side windows).

In another more advanced version of the system, the holographic projection system using micro-mechanical techniques may project the information directly upon the user's retina. This technique is detailed in co-pending patent application entitled "Remote Retinal Imaging Projection System" The variety of other potential applications of the present technologies are also conceivable which include:

# SYSTEM FOR PROVIDING REMOTE ACCESS AND MODIFICATION CAPABILITY OF DOCUMENTS AND E-MAIL BY A SENDER

The present system involves the use of a technique which enables the sender of a document, upon permission of the recipient, to provide subsequent updates and revisions to that document remotely and automatically in the absence of the recipient. Utilizing the computer's modem and e-mail delivery system, these updates could be provided remotely by the original sender (who alone can gain appropriate access privileges through a unique access code or authentication). The sender may also be able to identify whether or not the old version had been accessed by the recipient and possibly even at the level of the particular segment, or segments in which the specific changes had been made. (This capability could also be two way in which certification of not only receipt, but also access particular portions of the information by the intended recipient could be positively confirmed by the sender however, subject to approval by the recipient). If the recipient is amenable, such receipt confirmation could also be automatically time stamped by a secure trusted agent on the recipient's machine. The system could also perform the appropriate modifications in different drives on which the document is stored. If the



accessor's script could carry-out certified verification of more diverse or complex actions by a desired accessor, for example:

1. Verify (and if desired notify) the requester of the accessor having sent or received and read document (or e-mail) with content X to or from recipient or sender Y. For example, knowing that recipient had read a message or document sent by the sender (or another sender) or otherwise accessed may prompt requestor to take the next action or communicate with recipient.

Or verify (and if desired notify) the requestor (e.g., an employer) of the accessor (e.g., an employee) having made telephone contact with X individual(s) or phone number(s) at a certain time(s) (or within a certain time frame and/or containing certain message-type or content). Or the content (contextual profile) of the conversation(s) may be revealed or confirmed as being within a target contextual domain(s).

2. Verify (and if desired, notify) the requestor (e.g., an employer) of the accessor (e.g., an employee) having performed certain definable and/or desirable volume of on-line tasks, e.g., typed certain type(s) of content, performed certain accounting, billing (telephone handling courteous) prompt or other administrative functions. Knowledge by requestor (e.g., via notification) may then prompt requestor to take a subsequent action (or e.g., communicate again with accessor).

3. For any of the above notify requestor (e.g., employer) of the failure of the accessor to perform certain desired (and/or volume) of tasks with pre-defined threshold parameters.

4. Verify and, if desired, notify requestor (e.g., teacher or parent) of accessor (e.g., student or child respectively) of certain on-line content consumed, e.g., Web pages, paragraph by paragraph content (including wireless e-books) or on-line homework, such as quizzes, essays, reports, in conjunction with or independent from such on-line content.



1 5. Verify and, if desired, notify requestor (e.g., employer or teacher/parent) of accessor  
2 (employee or student/child respectively) of accessing, generating or conversing  
3 (acoustically/telephonically or textually) certain content which falls outside of a  
4 certain defined set of content clusters which are deemed "appropriate" within certain  
5 time periods (e.g., work-time such as when the user is "clocked-in", homework time,  
6 etc.) or in general. And/or verify or notify accessor of certain clusters, which are  
7 accessed or generated within certain content clusters. The above cluster based  
8 identification rules may also be definable by not only time of day or day of week, but  
9 domain, e.g., e-mails and phone conversations of a "social" context or style are  
10 inappropriate while pornography is inappropriate during work time for employees or  
11 anytime for students or children. The present "conversation" based variation, may  
12 even be implemented in conjunction with a passively collected acoustic monitoring  
13 via the client's microphone of an employee's, student's or child's spoken  
14 conversations in/around his/her (or a) client machine. Identity may be determined in  
15 conjunction with associated voice print information and bias results to documents,  
16 which match that style in accordance with the appropriate weighting criterion  
17 matching the user profile.

18  
19 b. Give the user a set of selection criteria to bias the search (or recommended  
20 navigational links) by those styles, which s/he desires presently.

21  
22 In addition, the present system could further actually use style as an additional means for  
23 identifying documents, which may be relevant to a search, e.g., determine what style(s) a  
24 preponderance of document in the present search tend to possess, then add this criteria as  
25 an additional weighting criteria to determining document relevance (or finding new  
26 potentially relevant documents). I.e., as part of the information retrieval (attribute  
27 cluster) model, use "style" as an additional attribute for and similarity to other documents  
28 and queries.

29  
30 System for Mining User Reactions and Responses to On-line Media by User Profile  
31 Characteristics:

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31

As described, are some primary examples of how the present technique can usefully be deployed. User profiling may also be usefully applied for purposes of collecting feedback about various types of user reactions and responses to various on-line stimuli (e.g. general Web and e-commerce sites to various products such as movies, music, interactive content, advertising news, interactions with other users etc. A somewhat related methodology was disclosed in the parent patent involving a technique for allowing users to rate vendors according to a variety of criteria and enabling the vendors (e.g. for marketing purposes) and/or users to observe statistical correlations between the user profile characteristics and ratings (using data mining techniques). In a simple example, application of the present methodology, the users viewing a Web page (containing any of the above suggested informational types) are allowed to provide their personal comments about the page or about their own experience in viewing the page. Informational retrieval and statistical NLP techniques may be used to cluster (using standard clustering techniques) both the comments (by similarity of their content profiles) as well as the users (by similarity of their user profiles). In the first instance, the comment cluster exemplars are extracted, used as the comments which most closely exemplify each of their associated clusters and thus presented to vendors and/or users. The aggregate (average) of the user profiles associated with that cluster, the user profile of the user who provided the exemplar comment or the subset of user profiles which characterize subclusters of the comment clusters (wherein the subclustering routine is based upon the profiles of users within each comment cluster) are also presented to the users or associated vendors in conjunction with the exemplary comment most closely associated with that stereotypical profile of users. In a variation, which could be an automatically selected alternative if the exemplar fails to provide a clear representation of the various comments in the cluster (if further subclustering does not adequately achieve this objective), it may be possible to “combine” these various comments which are closest to the clusters’ centroid by identifying those portions of the comments which are similar in meaning and those portions which are different. For similar portions, by preferentially utilizing the counterparts which are closest to the centroid and adding to it those portions which are different in piecemeal fashion, the desired objective may be achieved.

0069098.102700

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30

In a variation, the different types of comments may be clustered and accessible to users via a hierarchical cluster tree used to create a menu of automatically labelled clusters (see parent patent detailing this general method). Users may access such menus by combining querying (of user or content attribute with menu navigation).

It is of course reasonable to combine the technique in the parent patent application for statistically correlating user profile attributes with ratings associated with the users possessing those attributes.

It is also reasonable to apply the present technique to video (or streaming video content) whereby during the course of viewing the content relevant survey questions are presented to users and if desired, the correlation statistics of the results with the associated user attributes are presented to the content owner and/or (preferably) also future viewers. In a novel and advanced variation, time shifted video technology would enable the ability to capture audio or full video/audio versions of the comments of the user.

In a very novel application, it is possible upon permission of the associated users, to automatically construct menu trees (in accordance with the above method as disclosed in the parent case) of live voice or text chat in real-time dynamic fashion. Spoken conversations may be automatically profiled using acoustic speech-to-text methods. In accordance with the methods for automatic construction of virtual communities (in this case for dynamic chat) as disclosed in the parent case, these communities can be constructed automatically and on-the-fly based upon similarities of chat content, queries, navigational (content) selection from automatically constructed menu and/or user profile similarity or attributes of the profiles of the users which may be selected. Thus, in the former case a user could in the case of a pull down menu gain access by title of key phrase (which are ideally automatically and dynamically created and updated in accordance with the content of dynamic spoken or recorded conversation spoken or typed in real time and which may be further dynamically categorized and re-categorized within a hierarchical menu structure format. Although, the



1 word(s)/phrase(s)(again, so long as the user's profile is within the disclosure policy of the  
2 transmitting user. If a willing recipient so desires for both the chat and instant messaging  
3 variations, s/he may receive a copy of the text or voice transmission or if the user is not able  
4 to access it in real-time or otherwise. Conversely, the user may as a sender wish to sent the  
5 transmission to

6 1. All users who have indicated potential interest in the contents (and/or profile or identity  
7 of the sender) or users who match a particular user profile (e.g., interest or proficiency such  
8 as which is of measured similarity to the contents of the transmission) and/or

9  
10 3. Users who are in the process of engaging in a user-user (or multi-user) dialog, which is of  
11 measured similarity to that of the transmission.

12  
13 An example application of the present system includes, for TV viewers, the ability to gain  
14 access to different types of feedback from other users who are also presently viewing or had  
15 previously viewed a particular TV program, e. g., a comedy, a news story or political speech  
16 in which the present technique could be used effectively as a filter allowing certain types of  
17 user or user attributes to reach the user and others to be suppressed or squelched. In the  
18 latter example, a user who considers herself to be a female liberal may provide settings  
19 during a campaign speech by a conservative Republican speaking out against abortion to  
20 listen to acoustic/verbal gestures by other liberal females advocating abortion). Audible  
21 comments may be unfiltered and heard by the user, if desired.

22 If the program is a re-broadcast, more elaborate filtering capabilities are conceivable  
23 involving statistical analysis of the spoken language content (such could be conceivably  
24 performed for real time live information, however, at the expense of a slight delay). As  
25 suggested, it may be also possible to observe textually (e.g., through speech to text  
26 techniques) or hear spoken conversations as they exist between individuals possessing  
27 desired attributes, containing content characterized by key words or phrases from a menu  
28 selection or provided by/between a desired individual or individuals respectively.

29  
30 In another application, users viewing video or streaming video content may upon their  
31 permission, agree to disclose their user profiles and be acoustically monitored such that

1 comments, verbal, verbal/audible gestures and expressions and/or video of Effectuated thereby  
2 may be of interest to users

3  
4 In conjunction with an educational program a user may wish to observe comments by  
5 those other individuals who are extremely knowledgeable in the field (to collect  
6 supplemental information and other feedback about the presently viewed content).  
7 Similarly, in the political speech example the user may desire to hear the gestures and/or  
8 spoken comments, reactions and/or brief dialogues occurring during the course of the  
9 political speech by individuals sharing the same political views as the user and which are  
10 considered to be part of the educated elite. In another variation, individuals who  
11 represent the exemplars of different cluster of users (and/or those which are manually  
12 selected as providing interesting and/or entertaining feedback) could be heard by the user  
13 collectively and the reactions expressed as acoustic feedback could be statistically  
14 analyzed in dynamic fashion and presented to the user as a breakdown of the user  
15 attributes which presently characterize the present reaction or response of the virtual  
16 audience.

17  
18 In a future application, such system could be extended to such things as immersive virtual  
19 reality (stationary or non-stationary) systems and/or video gaming systems in which the  
20 reactions of characters (or even different "personas") to various situations and events  
21 could be developed from data collected from the reactions of actual users comprising the  
22 different user clusters as above suggested.

23  
24 Applications to Constructing User Profiles and Matching Users by Similarities in Their  
25 Social and Psychological Profiles and/or Life Circumstances and Experiences

26  
27 There is an untapped opportunity with potentially considerable and deep implications  
28 which could be used to enable very detailed assessment and associated profiling of  
29 individuals. The parent patent describes a very comprehensive methodology by which  
30 users could potentially be profiled so as to define their preferences across almost any kind  
31 of content or commercial products and services. Surveys used to reveal psychological

09699098-102700

1 traits are also used. In the present extension, we first suggest a primary system  
2 methodology in which extremely detailed information is collected and aggregated into  
3 the user profile particularly pertaining to the user” psychological profile (user  
4 preference/interest information demographics, etc. are also useful data in that there are  
5 likely correlations which exist with certain psychological attributes of the user). In the  
6 clinical field of psychoanalysis, much information regarding the psychological and  
7 psycho-pathological characteristics of the user is determined, however in compliance  
8 with the confidentiality requirements of the doctor-patient privilege, most if not all of this  
9 information remains isolated and never accessible by a large-scale statistical database for  
10 purposes of cross-correlation of psychological characteristics, phenomena, psycho-  
11 pathology (as well as other clinical pathology), as well as behavior, interests, preferences  
12 and more specifically identifiable behavior patterns as well as notable life experiences.  
13 Part of the answer to this dilemma may well lie in the application of secure data  
14 disclosure to SDI, the pseudonym proxy server and (for extra security) randomized  
15 aggregates for purposes of harvesting the statistical value within the data and (if needed),  
16 interacting with the user while maintaining completely secure and private individual user  
17 data. The present system suggests that there may be much greater accuracy which could  
18 be achieved in the profiling of individuals for use in a number of applications such as the  
19 psychological modeling of patients for use in the diagnosis, treatment, prediction  
20 (including predictive treatment) of psycho-pathology, the matching of “similar” patients  
21 together (for group therapy and/or pseudonymous electronic or physical mail  
22 correspondences) for providing support, inter-patient communication, counseling, etc.  
23 (for which present traumatic life experiences, if relevant, may be quite applicable as  
24 well), as well as the extended/improved modeling/prediction of user interests/preferences  
25 enabled by SDI.  
26 Another correlated data type could be employee file databases through which correlations  
27 and predictions of employee behavior and performance, recommendation as to remedies  
28 for existing problems and performance enhancing tactics again could be achieved.  
29 Unfortunately not all data is likely to be accessible for all users (e.g.,  
30 browsing/transaction histories are likely to be available but often not information  
31 regarding a user’s life experiences or psychological history). To fill in this missing data,





1 example, disclose to the employer "Mr..X has a psychological profile feature which is  
2 highly correlated with a high aptitude for the present job which s/he is applying for and  
3 past job experience Y is 50% more likely than the average employee for that position to  
4 lead to a subsequent promotion and several "employee of the month awards"

5  
6 Application to Identifying Similar Individuals Behaving Under Similar Circumstance  
7

8 The present user profile information could also be extremely relevant for a novel Internet  
9 based service by which individuals could pseudonymously release their profiles and/or  
10 define specific situations and/or conditions (which could be actual, hypothetical or  
11 combination thereof) and the system will identify other scenarios (of like pseudonymous  
12 individuals and/or circumstances) to the general scenario and emphasizing and  
13 prioritizing certain particular aspects thereof. Examples include identifying individuals  
14 who had previously (or are presently) in the process of making a very similar professional  
15 decision, interpersonal decision.(such as regarding a present or a potential spouse).  
16 Ideally, such interpersonal decisions would involve a very similar situation (such as a  
17 decision), a very similar individual in the analogous position of the user (and ideally, if  
18 possible other similar party(s) fulfilling the analogous role(s) as the present situation  
19 which confronts the user. The relative importance (priority) of the degree of similarity of  
20 various analogous roles in an analogous situation is determined in part by the type of  
21 information which the user would like to determine. For example, determining what a  
22 particular individual would do in a similar circumstance may depend upon how  
23 dependent that user's action .or decision is upon the circumstances and conditions versus  
24 another individual(s) (and which individual(s) and/or the two (or more) individuals  
25 involved to which the context of the situation primarily relates. Of course, the nature of  
26 the action in question may predicate the relative importance of one or more of the above  
27 factors. These determinations are often quite subjective in nature. Requestees may also  
28 (if the choice exists) be situations which are either of a historical nature (i.e., the situation  
29 had occurred in the past) or of a present nature. In some cases in the latter case, is  
30 preferable as the user may actively request, e.g., the party to which s/he directly relates or  
31 is potentially affected by to provide advice and/or feedback or to what s/he would do



1 within the section entitled "Resolution credentials" in the access reachability introduction  
2 and profile disclosure between two or more parties as disclosed

3  
4 In order to utilize LEIA as a user agent, which is integral and ubiquitous within the work  
5 environment, it is necessary to make its learning as seamless and invisible from the user's  
6 perspective as possible.

7  
8 The next level in applying the capabilities of LEIA is in being able to implicitly learn the  
9 context of the user's present temporal activities as well as establish relative priorities of  
10 the activities that the user (and others who intend to interface with the user) are engaged  
11 in and from the user's previous behavior to similar conditions automatically generate  
12 appropriateness functions (or rules) to automate the process of handling meeting  
13 schedules and filtering real time requests for correspondence with the user. It is possible  
14 to infer:

- 15  
16 1. The content profiles which are associated with the user's present activities.  
17  
18 2. (Often obviously) The nature of the user's present activities by observing what sorts of  
19 actions the user is engaged in, in the office.  
20

21 These content profiles and activity indicators may be passively observed by the user's  
22 present interactions (and timing thereof) with his/her PC, other smart appliances and data  
23 and voice communications which the user engages in (telephonically or from standard  
24 face to face dialogues) e.g., suggesting when the user is reading the morning news,  
25 checking phone messages, responding to e-mail, engaged in a particular project, etc. If  
26 these clues are not observable (or not positively identifiable) cyclical time dependent  
27 patterns may be used as implicit indicators of the user's present activities (see the above  
28 section entitled "Further Applications of LEIA"). Alternatively, the system could prompt  
29 a request of confirmation of the user's present activities and possibly on clues such as  
30 verbal regarding present and future established or changing activities and plans These  
31 same methods for utilizing time pattern elements which are used to predict the location of

1 the user may be readily extended to predict or help to substantiate other statistical user  
2 variables relating to the types of actions and content related profiles characterizing the  
3 user's temporal behavior patterns. The following application of these techniques is used  
4 to enhance the automatic meeting scheduler:

5  
6 The identity of the task or request, the activity and content attributes of the task or request  
7 profiles (target object profiles) may be developed utilizing the priority of that task  
8 relative to the user. This priority is estimated from the priority of the task/request (and its  
9 associated attributes) over other tasks/requests by the user. In order to better substantiate  
10 the relative importance to the user as inferred regarding these activities, the system  
11 should also display the scalar ratings which the user could adjust manually.

12 Nonetheless, there may be a degree of uncertainty in the user's intentions and LEIA's  
13 inference of the user's intentions. Moreover, unanticipated changeability by the user  
14 (uncertainty) may occur perhaps with increased frequency under certain variables like the  
15 identity of the requestor, the time of day, context of the users other present activities or  
16 when other meetings and obligations are pending which possess a relatively greater  
17 priority with the user (i.e., if they often occur spontaneously). In order to take this  
18 uncertainty factor into consideration, an overall statistical estimate may then be  
19 calculated taking into account the predicted statistical probability of each party ultimately  
20 being available for the meeting. This serves the purpose of both refining the automatic  
21 scheduling process and informing would be visitors before hand as to the relative  
22 importance that their prospective meeting is (or increasingly becomes as the meeting time  
23 approaches) to the employee as well as the overall estimated degree of certainty of the  
24 meeting. This feature may be particularly useful for impromptu visitors where  
25 considerable uncertainty exists as to whether and to what degree the visitor is imposing  
26 upon the employee's schedule.

27  
28 The temporal context of user behavior (using temporal time series analysis in conjunction  
29 with use of content analysis and user behavior assessment may be both important  
30 predictors as to whether a prospective meeting of a disclosed purpose is worthwhile for  
31 the employee based on present/future priorities and at which most likely times. By far,

1 its most valuable application , however is in its ability to perform content analysis and  
2 user behavior assessment in real time in order to enhance LEIA's ability to determine  
3 whether and to what degree a would-be unannounced visitor or telephone or intercom  
4 callers unannounced contact is relevant to the employee" present activities and if not to  
5 coordinate with the visitor's agent the time or times when it would be most relevant to the  
6 employees anticipated activities for example there are certain activities during which an  
7 employee doesn't wish to be usually disturbed, e.g. while reading the morning news,  
8 taking a coffee break, calling home, responding to an important email message, ten  
9 minutes before an important meeting. Often within an organization individuals will  
10 attempt to speak with one another as such , the intended duration affects priority as does  
11 the identity of the individual in the organization. Often this variable can be automatically  
12 predicted by LEIA . In such a situation, the relative priority of the user's preexisting  
13 priority tasks is weighted against the priority of the requested tasks prior to notification.  
14 The above application can be readily extended as well to phone calls, in which the user  
15 could be automatically prompted by LEIA to disclose his/her purpose of the call and/or  
16 his/her identity if necessary. This application would be a useful extension to the  
17 telephony variation of the email filter described in issued patent entitled "System and  
18 Method for Electronic Identification of Desirable Objects in which the above additional  
19 inputs about employees and visitors may provide additional useful feedback in  
20 automatically determining the rules dictating how to best handle the call or forward it to a  
21 more appropriate employee on behalf of the employee. If there is a certain degree of  
22 uncertainty in the agent's priority determination or assessment of the task affected the  
23 most appropriate rule the employee may observe the rule just prior to automatic  
24 implementation. Or the action which it represents. The present value of the contextual  
25 employee activity assessment method as disclosed, if extended to advertising could be  
26 significant. For example the knowledge of the context of the user's present activities  
27 could determine what type and when an advertisement targeted to an employee, e.g. a  
28 decision maker will most likely elicit the best response (where the advertisement  
29 matches the professional interest profile) is an ad which corresponds with particular  
30 activities of the user ( e.g., which relate in some way with the ad/promotion. As  
31 discussed direct voice telephony and email dialogues may provide very useful clues

